

Note: this is the authors' final draft; all citations should refer to the final published version, in the *Journal of Medicine and Philosophy*; see <https://academic.oup.com/jmp>

Reconciling regulation with scientific autonomy in dual-use research

Nicholas G. Evans, Michael J. Selgelid, and Robert Mark Simpson

Abstract. In debates over the regulation of communication related to dual-use research, we have to weigh the risks that such communication creates, against the value of scientific autonomy. The censorship of such communication seems justifiable in certain cases, given the potentially catastrophic applications of some dual-use research. But this conclusion gives rise to another kind of danger: that regulators will use overly simplistic cost-benefit reasoning to rationalise excessive regulation of scientific research. In response to this, we show how institutional design principles and normative frameworks from free speech theory can be used to help extend the argument for regulating dangerous dual-use research beyond overly simplistic cost-benefit reasoning, but without reverting to an implausibly absolutist view of scientific autonomy.

1. The ethical challenges of dual-use research

Some scientific research has the potential to be used for both very beneficial ends and very harmful ends. For example, research that can be used to develop a new life-saving vaccine might also enable a malevolent actor to engineer a pathogen that could be used to carry out acts of bioterrorism.¹ The governance of this so-called 'dual-use research' has become a hotly debated policy issue in the life sciences. The challenge for policy-makers is to manage the risks that are posed by the misuse of scientific and technological

¹ Seumas Miller and Michael Selgelid, *Ethical and Philosophical Considerations of the Dual-Use Dilemma in the Biological Sciences* (Dordrecht: Springer, 2008).

research, but without unduly compromising the autonomy of scientists or forgoing the benefits of their research.

As we will argue below, the magnitude of the risks in this area mean that some government censorship of dual-use research is necessary and justifiable. But this kind of reasoning is dangerous in its own way. Repressive governments often rationalise their censorship of disapproved speech by adverting to the risks and dangers that it poses. An overly simplistic cost-benefit approach to debates around censorship runs the risk of devolving into an alibi for authoritarianism. Our broader aim, then, is to show how the case for regulating dangerous dual-use research can be augmented, using institutional principles and normative frameworks from free speech discourse, in a way that addresses these kinds of concerns, but without reverting to an implausibly absolutist view about the value of scientific autonomy.

To set the background for this discussion, we begin with a brief overview of the evolving policy debates and conflicts surrounding dual-use research in the U.S. over the last decade or so. Despite acknowledging the risks of dual-use research, the scientific community has generally opposed measures that would allow government agencies to heavily regulate the dissemination of dual-use research. The U.S. National Research Council (NRC) acknowledged the *prima facie* case for government censorship of dual-use research as early as 2004, but it argued that this would unduly impede the progress of biological research and the beneficial applications of novel scientific discoveries, and that policymakers should instead rely on the scientific community's voluntary self-governance (in matters pertaining to dissemination of information).² The problem with this approach is that scientists and other members of the scientific community (e.g. publishers) generally are not well-placed to accurately assess the risks related to dual-use research, given that these risks often relate to national security issues that fall outside the scientist's sphere of expertise and/or because adequate risk assessment may sometimes require access to classified information (unavailable to ordinary scientists).³ Moreover, given the importance of publication to the scientist's career advancement, and the publisher's profits, a regime of voluntary self-regulation will sometimes result in the publication of papers that pose serious security dangers. If members of the scientific community are free to make their own decisions about such matters, we can expect that they will sometimes decide to allow the dissemination of research despite its potentially catastrophic dangers.

² National Research Council (NRC), *Biotechnology Research in an Age of Terrorism* (Washington DC: National Academies Press, 2004). This view was later endorsed by the NRC's follow up report, *Globalization, Biosecurity, and the Future of the Life Sciences* (Washington DC: National Academies Press, 2006).

³ Michael Selgelid, "A tale of two studies: ethics, bioterrorism, and the censorship of science", *The Hastings Center Report* 37/3 (2007): 35-43.

This is not just conjecture. The U.S. National Science Advisory Biosecurity Board (NSABB) has backed up the NRC’s position: although the dissemination of dual-use research can create serious risks, according to NSABB, government regulators should not override the views of scientists and publishers about when dissemination is appropriate.⁴ Shortly after its formation in 2005 the NSABB was consulted about a paper accepted for publication in *Science*, which described the reconstruction of the 1918 “Spanish flu” virus that killed 50-100 million people.⁵ The NSABB board unanimously decided that the paper should be published, despite the associated risks. What’s more, *Science*’s Editor-in-Chief, Donald Kennedy, said that he would “have published the paper *even if* the NSABB had voted otherwise”.⁶ In this case, then, an influential actor in the scientific community was prepared to publish a paper even in the face of an assessment, by a relevant body of experts, that this would gravely threaten human health and security.

The debate around this issue came to a head in 2011, when two separate teams of scientists created genetically modified strains of highly pathogenic avian influenza H5N1 that were airborne transmissible between ferrets (which provide the best model for influenza in humans). Ordinary H5N1 only infects humans via birds, and it does so with little efficacy. As of 2016 there have only been 854 recorded cases of H5N1 in humans, but the fatality rate for those infected is 52%. The release of a similarly deadly, human-to-human transmissible, airborne strain of H5N1 could thus be catastrophic. To illustrate the danger, recall that the 1918 Spanish flu pandemic had a fatality rate of only 2.5%.⁷ An influenza pandemic resulting from airborne, human-to-human transmissible H5N1, if its fatality rate were comparable to ordinary H5N1, could result in between 20 million and 1.6 *billion* deaths.⁸

Concerns about these dangers and the potential misuse of the H5N1 studies prompted NSABB to recommend that key details of the methods by which the modified viruses were produced be omitted

⁴ National Science Advisory Board for Biosecurity, *Proposed Framework for the Oversight of Dual-use Life Science Research: Strategies for Minimizing the Potential Misuse of Research Information* (Washington DC: National Science Advisory Board for Biosecurity, 2007).

⁵ Jeffery K. Taubenberger and David M. Morens, “1918 influenza: the mother of all pandemics”, *Emerging Infectious Diseases* 12/1 (2006): 15-22.

⁶ Donald Kennedy, “Better never than late”, *Science* 310/5746 (2005): 195 [our emphasis].

⁷ Taubenberger and Morens, “1918 influenza”.

⁸ Note, however, that the airborne strains created via this research weren’t so deadly – for ferrets, anyway. See Marc Lipsitch and Thomas V. Inglesby, “Moratorium on research intended to create novel potential pandemic pathogens”, *mBio* 5/6 (2014): e02366-14; Marc Lipsitch and Thomas V. Inglesby, “Erratum for Lipsitch and Inglesby, Moratorium on research intended to create novel potential pandemic pathogens”, *mBio* 6/1 (2015): e02534-14.

from articles (under review in *Science* and *Nature*) reporting findings of the studies in question.⁹ Although NSABB eventually reversed this recommendation in 2012 in response to revisions to the two papers, and further (purported) justification of the public health benefits of the studies, its earlier recommendations lacked any binding force, given its merely advisory powers.¹⁰ Regardless of whether NSABB changed its position, final decisions about what details to publish ultimately rested with the editors and authors.

After the controversy surrounding this case, a 2012 policy released by the U.S. Government stated that “if the risks posed by the research cannot be adequately mitigated [through other measures] Federal departments and agencies will determine whether it is appropriate to... request voluntary redaction of the research publications or communications [or to] classify the research ... in accordance with National Security Decision Directive NSDD-189”.¹¹ This guidance seemed to allow funding agencies and departments to classify or terminate H5N1 research projects with risks that couldn’t be mitigated. However, a subsequent report to the U.S. National Academy of Sciences Committee on Science Technology and the Law has noted that Directive NSDD-189 protects “fundamental research” from classification and restrictions under export control laws, to the maximum extent possible. This “fundamental research” standard has historically encompassed any life sciences research conducted in unclassified laboratories, for the purpose of publication in academic journals. According to the report, this standard makes it “exceedingly difficult for these [regulatory] regimes to effectively control dissemination of [dual-use research] if the researchers intend to publish their research results... and have no funding-related restrictions on their freedom to do so”.¹²

A similar set of biosafety and biosecurity concerns has arisen around so-called “gain-of-function” (GOF) research, i.e. research that leads to the creation of pathogens, including the modified H5N1 varieties discussed above, that are more virulent or transmissible than their natural counterparts. In 2014, in response to these growing concerns, the U.S. Government imposed a pause on public funding for GOF studies – involving influenza, severe acute respiratory syndrome coronavirus (SARS), and Middle East

⁹ National Science Advisory Board for Biosecurity, *National Science Advisory Board for Biosecurity Recommendations* (November 21, 2011).

¹⁰ National Science Advisory Board for Biosecurity, *Findings and Recommendations* (March 29-30, 2012).

¹¹ Office of Science Policy, *United States Government Policy for Oversight of Life Sciences Dual Use Research of Concern* (Bethesda, MD: Office of Science Policy, 2012).

¹² Doron Hindin, Kim Strosnider, and Peter D. Trooboff, “The role of export controls in regulating dual use research of concern: striking a balance between freedom of fundamental research and national security (January 20, 2017); retrieved June 5, 2017 from http://sites.nationalacademies.org/cs/groups/pgasite/documents/webpage/pga_176436.pdf.

respiratory coronavirus (MERS) – and initiated a ‘deliberative process’ to assess the risks and benefits of such research.¹³ The worries around dual-use research like the H5N1 studies had expanded beyond the potential misuse of the research findings, to incorporate concerns about biosafety (i.e., dangers associated with laboratory accidents) .¹⁴ The NSABB’s report on these studies concluded that some life sciences research, including some GOF research, should not be conducted at all, because the risks of accidental or intentional release of novel pathogens outweighed the potential benefits of the research.¹⁵ While it recommended the careful communication of dual-use results, it remained silent on the issue of whether censorship was ever justified, except to note that

The U.S. government has no authority to mandate redaction, restriction, or classification of a scientific publication that it does not own or control, and the development of a mechanism for restricting communication of unclassified information to only those who require access, remain challenging and to date unsuccessful.¹⁶

The 2017 policy recommendations that emerged from the process initiated in 2014 made mention of the censorship of dual-use research in the context of requesting voluntary redaction of publications or communications resulting from potentially risky projects.¹⁷ It also advised federal agencies to make “classification determinations within the scope of their classification authorities and appropriate classification guidelines or ... consult with other agencies to make these determinations”.¹⁸ But without further clarification, this guidance is presumably subject to the strict constraints of NSDD-189 described above. In short, despite the legitimate views of some policymakers – that certain forms of dual-use research shouldn’t be

¹³ White House Office of Science and Technology Policy, *U.S. Government Gain-of-Function Deliberative Process and Research Funding Pause on Selected Gain-of-Function Research Involving Influenza, MERS, and SARS Viruses: Frequently Asked Questions* (Washington DC: White House Office of Science and Technology Policy, 2014); retrieved September 14, 2017 from <https://www.phe.gov/s3/dualuse/Documents/gof-qanda.pdf>.

¹⁴ Nicholas G. Evans, Marc Lipsitch, and Meira Levinson, “The ethics of biosafety considerations in gain-of-function research resulting in the creation of potential pandemic pathogens”, *Journal of Medical Ethics* 41/11 (2015): 901-8.

¹⁵ National Science Advisory Board for Biosecurity, *Recommendations for the Evaluation and Oversight of Proposed Gain-of-Function Research* (May 6, 2016): 4.

¹⁶ *Ibid.* 33.

¹⁷ Office of Science and Technology Policy, *Recommended Policy Guidance for Departmental Development of Review Mechanisms for Potential Pandemic Pathogen Care and Oversight (P3CO)* (January 9, 2017).

¹⁸ *Ibid.* 5.

openly communicated, and that some should not even be *conducted* – there has been little principled guidance, to date, about when dual-use research in the life sciences can be permissibly censored, including, *inter alia*, through government classification.

2. Regulating dual-use research: the current debate

In the sparse, but growing, ethical literature on the regulation of dual-use research, the arguments against censorship commonly appeal to the overriding importance of maintaining autonomy in scientific practice. One version of the argument stresses the benefits of scientific autonomy in promoting scientific knowledge as an end in itself. The advancement of scientific knowledge requires conditions of openness, where scientists are free to design, perform, and publish their research based on their own judgements and discretion.¹⁹ External regulation hampers this by increasing redundancy, diverting researchers from the most promising lines of inquiry, and impairing their ability to build on the work of their peers.²⁰ We should oppose regulation, then, so the argument goes, irrespective of the risks posed by dual-use research.

Another version of this argument appeals to the instrumental utility of scientific knowledge in promoting other valuable outcomes, in particular, improved human health and wellbeing. External regulation that restricts scientific communication impairs the realisation of these benefits. The fact that a scientific discovery could inadvertently result in the development of weaponised bioagents does not provide sufficient grounds for gagging researchers, so the opponents of regulation argue, because the benefits ultimately outweigh those risks.²¹ Relatedly, David Resnik argues that even if censorious regulatory powers are never actually applied, the mere threat of censorship could have a chilling effect on scientific research, and therefore, by extension, could jeopardise the beneficial results of scientific progress.²²

¹⁹ This kind of justification for scientific autonomy appears as far back as Michael Polyani, *The Logic of Liberty: Reflections and Rejoinders* (London: Routledge and Kegan Paul, 1951): 34.

²⁰ David Resnik, “Scientific autonomy and public oversight”, *Episteme* 5/2 (2008), 220-38; see also National Academies of Science, *Scientific Communication and National Security* (Washington DC: National Academies Press, 1982).

²¹ Robert Carlson, “The pace and proliferation of biological technologies”, *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science* 1/3 (2003): 203-14; Michael Selgelid, “A tale of two studies: ethics, bioterrorism, and the censorship of science”, *The Hastings Center Report* 37/3 (2007): 35-43.

²² David Resnik, “Openness versus secrecy in scientific research”, *Episteme* 2/3 (2006): 135-47.

However, neither of these arguments provides a convincing in-principle reason to oppose government regulation of dual-use research. Given the magnitude of the risks involved, it is implausible to assume that open scientific research will always have net benefits, either in terms of the intrinsic value of scientific knowledge, or in terms of the wider consequences of scientific progress for society. The real-world cases discussed above already illustrate how the costs of openness and voluntarily self-regulation in this domain could dramatically outweigh the benefits. Of course there is room for debate about exactly how the H5N1 ferret studies should have been handled.²³ Regardless of the particulars in that case, however, it is easy to imagine hypothetical cases in which the expected costs of unregulated scientific autonomy are clearly in the negative. Consider a disease like smallpox, for example, which is believed to have killed more people than any other infectious disease in history, and three times more people in the twentieth century than were killed by all the wars of that period.²⁴ If a laboratory discovered an easy way to synthesise a pathogen as contagious, deadly, and untreatable as smallpox, the dissemination of their production methods could obviously be catastrophic.

In cases where there is a non-negligible likelihood of a catastrophic outcome resulting from dual-use research, it is reasonable to think that governments could, in principle, justifiably impose coercive regulatory constraints that limit the dissemination of the relevant information, e.g. by limiting communication to those who strictly need to be informed in order to develop protective measures against the threat.²⁵ Given this framing of the issue, it seems perverse to insist that scientists or publishers should be accorded complete discretionary powers in deciding when and how to disseminate such information. Even if scientific knowledge is valuable for its own sake, and is usually likely to have other societal benefits, there is no good reason to believe that its value and the value of its societal benefits will *always outweigh* whatever harms might result from such knowledge. The only way to defend this claim is to subscribe to an implausibly extreme, absolutist view about the importance of scientific autonomy.

²³ E.g. see Ron A. M. Fouchier, Reply to “Comments on Fouchier’s calculation of risk and elapsed time for escape of a laboratory-acquired infection from his laboratory”, *mBio* 6, e00407–15 (2015), doi:10.1128/mBio.00407-15.

²⁴ Michael B. A. Oldstone, *Viruses, Plagues, and History* (New York: Oxford University Press, 1998).

²⁵ This proposal is discussed in Selgelid, “A tale of two studies”.

3. Regulating dual-use research: insights from free speech theory

Even with this conclusion in hand, questions remain about how the decision-making processes around these issues should be approached.²⁶ The main concern we want to address here is how to make our framework sensitive to the real and significant dangers of dual-use research, but also *equally* sensitive to the *countervailing* dangers that come with assigning broad powers of censorship to government agencies. We should resist absolutism about scientific autonomy, because sometimes the risks of disseminating dual-use research are too grave; but, in order to guard against the second set of dangers, we must also resist overly simplistic cost-benefit approaches to regulatory decision-making. As noted above, one way to get a grip on the problems with a cost-benefit approach is to reflect on the fact that governments that persecute their ideological opponents routinely invoke claims about the dangers of the speech that they are suppressing in order to defend their actions. State censorship on the basis of claims about the dangers of disapproved speech – which is essentially a crude form of cost-benefit justification – is a standard move in the authoritarian playbook.²⁷

These are the kinds of concerns that underpin the general project of articulating and justifying principles of free speech.²⁸ Now, just about everyone, including strident free speech advocates, will agree in principle that dangerous communication sometimes really *does* need to be restricted in order to protect

²⁶ To date there is only a small body of philosophically-informed literature on the ethics of dual-use research that aims to address these questions, e.g. Miller and Selgelid, *Ethical and Philosophical Considerations of the Dual-Use Dilemma in the Biological Sciences*; Selgelid, “A tale of two studies”; Resnik, “Openness versus secrecy in scientific research”, Nicholas G. Evans, “Great expectations – ethics, avian flu and the value of progress”, *Journal of Medical Ethics* 39/4 (2013): 209-13; David B. Resnik, “H5N1 avian flu research and the ethics of knowledge”, *Hastings Center Report* 43/2 (2013): 22-33.

²⁷ For discussion of how claims about dangerous speech are used to defend oppressive ideological censorship in authoritarian regimes, see for instance Tomas Venclova, “USSR: stages of censorship”, *Index on Censorship* 7/4 (1978): 61-62, J. M. Coetzee, “The work of the censor: censorship in South Africa” in *Giving Offense: Essays on Censorship* (Chicago: University of Chicago Press, 1996): 185-203.

²⁸ The right to freedom of speech has been defended via multiple lines of argument, e.g. (i) by appeal to the utilitarian benefits of free speech (J. S. Mill, *On Liberty* (Amherst: Prometheus Books, 1986 [1859])), (ii) by appeal to the value of self-governance (James Griffin, *On Human Rights* (New York: Oxford University Press, 2008)), (iii) by appeal to the value of individual autonomy (for critical discussion see Susan J. Brison, “The autonomy defense of free speech”, *Ethics* 108/2 (1998): 312-39), (iv) by appeal to the demands of egalitarian respect (Kenneth L. Karst, “Equality as a central principle in the First Amendment”, *The University of Chicago Law Review* 43/1 (1975): 20-68), and (v) as a component of the basic structure of liberal democracy (Frederick Schauer, *Free Speech: A Philosophical Enquiry* (London: Cambridge University Press, 1982)). Only a few authors to date explicitly acknowledge that free speech principles are relevant to our thinking about the regulation of dual-use research. Selgelid notes that government-backed censorship of dual-use research “may threaten freedom of speech more generally”; see “A tale of two studies”: 41. And Resnik says part of the reason to minimise restrictions on the publication of dual-use research is because they would infringe on people’s speech rights; see “Openness versus secrecy in scientific research”.

us against extraordinary risks.²⁹ In practice, however, if we allow government agencies to restrict communication that they deem to be unacceptably dangerous – where this verdict is purely based on their own cost-benefit assessment, without any additional systems of oversight and accountability – we are not taking the second set of dangers seriously. Authorities with strong regulatory powers over the communication of researchers can abuse their power, either due to self-serving motives, incompetence, or both.³⁰ Part of the idea of a *right* to free speech is to guard against this, by imposing constraints that limit the ability of authorities to pursue their aims through the regulation of communication. Recent cases in which notionally liberal governments have imposed heavy-handed and apparently ideologically motivated gags on scientists lend further support to these worries.³¹

Even in cases where regulatory authorities are suppressing dangerous communication in pursuit of legitimate aims, problems can still arise under a crude cost-benefit framework for decision-making. For instance, under this kind of framework, the agencies responsible for the regulation of scientific research will be subject to perverse incentives in their judgements. They will have an incentive to underestimate the probability-weighted benefit of potentially harmful research, because the benefits of that research will usually be remote, and credit for those benefits, if and when they're realised, will not typically be ascribed to the regulatory agency that opted to protect the research environment. Regulatory agencies will also have an incentive to *overestimate* the probability-weighted *harm* of potentially dangerous research, since the costs that will result if harmful outcomes eventuate are likely to be blamed on the regulatory body that

²⁹ For instance, after arguing for a categorical right to free speech, Nagel nevertheless concedes that “it is compatible with this conception of [free speech] rights that they are not absolute, and that there may be some threshold, defined in consequential... terms, at which they give way”; Thomas Nagel, “Personal rights and public space”, *Philosophy & Public Affairs* 24/2 (1995): 83-107, 100. More generally, contemporary theories of rights generally allow that even the most fundamental individual rights can be infringed upon in order to prevent especially grave harms; for discussion see e.g. Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974): 30.

³⁰ Indeed, claims about the inability of governments to regulate public discourse in a judicious and even-handed way represent one important strand of free speech theory. Frederick Schauer is a prominent representative of this view; see *Free Speech: A Philosophical Enquiry*; “Principles, institutions, and the First Amendment”, *Harvard Law Review* 112/1 (1998): 84-120; “First Amendment opportunism” in Lee C. Bollinger and Geoffrey R. Stone (Eds), *Eternally Vigilant: Free Speech in the Modern Era* (Chicago: University of Chicago Press, 2002).

³¹ For a useful overview of the gags and other regulatory controls that Canada's Harper government imposed on scientists whose work had implications in relation to climate change policy, see Sarah Zhang, “Looking back at Canada's political fights over science”, *The Atlantic* (January 26, 2017).

failed to prevent them.³² We can thus expect regulators to assess the expected costs and benefits of regulatory intervention in a way that does not accurately correspond with the costs and benefits for society at large, and in a way that – without mechanisms to counteract this – will tend towards the over-regulation of scientific communication.

There are also problems of a more theoretical nature that come with using a simplified cost-benefit framework to govern the regulation of dual-use research. A simplified cost-benefit defense of regulation has this basic form: “if the probability-weighted costs of leaving dual-use research unregulated outweigh the probability-weighted benefits of leaving dual-use research unregulated, then regulation is justified in principle”. The worry is that this kind of reasoning, on its face, will allow that *whenever* a regulatory agency deems that the expected utility of a scientific research program has a negative value – given the probability-weighted magnitude of its potential harms – censorship is justified. If agencies overseeing dual-use research apply this kind of adjudicative framework, then we would expect to see exactly the types of chilling effect that Resnik and others are worried about. Other things being equal, researchers will be averse to pursuing projects that they fear might be gagged, shut down, or otherwise inhibited by regulatory agencies. If they are working against the backdrop of a crude cost-benefit adjudicative framework, like the one we have sketched, then researchers will tend to refrain from pursuing not only the projects that they recognise as genuinely dangerous, but also projects for which they are uncertain about whether the relevant regulatory agency would deem it to be dangerous. This is precisely how chilling effects function. Uncertainty about whether the speaker’s speech will be regarded as a breach of the relevant norm encourages the speaker to censor herself, rather than taking on an unpredictable risk of incurring a significant penalty after the fact.³³

The natural response to this is to look for a way to modify the cost-benefit decision-making calculus so that researchers are not navigating in the dark, and so that it is harder for regulatory agencies to make implausible, idiosyncratic, or otherwise difficult-to-predict judgements about the dangers of research. This is partly a problem of institutional design, but we also need a normative framework to underwrite this theoretical modification.

³² Schauer’s work has been important in articulating how the structural form of state censorship leads to these kinds of distortions; see in particular “Slippery slopes”, *Harvard Law Review* 99/2 (1985): 361-83.

³³ For further analysis of how chilling effects work, see Leslie Kendrick, “Speech, intent, and the chilling effect”, *William and Mary Law Review* 54/5 (2012): 1633-91; Frederick Schauer, “Fear, risk and the first amendment: unravelling the chilling effect”, *Boston University Law Review* 58/5 (1978): 685-732.

This is less straightforward than it might seem. We cannot modify the cost-benefit calculus by giving researchers an absolute right to operate autonomously. That ignores the problem we started with: some dual-use research is too dangerous to be categorically protected against regulatory controls. Another approach would be to assign a special weight to scientific autonomy in our cost-benefit calculus. In some cases, where the expected costs of dual-use research seem to outweigh the expected benefits, we might nevertheless leave the research unregulated, based on the judgement that respect for scientific autonomy places a ‘thumb on the scales’ in favour of openness rather than regulation.³⁴ But if this weighting is based on the discretionary assessment of regulators trying to balance the costs and benefits against each other, this adjustment to our framework will do little to alleviate the uncertainties that create the chilling effects. Another alternative is to require compelling evidence on behalf of regulators before censorship can be enforced, and allow research to proceed until such time as that evidence can be established – placing an *epistemic* (rather than normative) thumb on the scales against censorship.³⁵ But this will just delay the potential chilling, rather than overcoming it. What we need is a normative framework that supports the consequentialist reasoning used to justify the suppression of dangerous research in *exceptional* cases, but which simultaneously offers assurance to researchers that – *in the normal range of cases* – their autonomy will be respected, and their projects will not be gagged or de-funded due to spurious, erroneous, idiosyncratic, or ideologically-motivated assessments by regulators about the dangers of their work.

One could try to downplay these concerns. Even if there are risks of excessive interference in this arena in theory, governments to date have not taken a heavy hand in regulating dual-use research, as noted above. But this is too complacent. The question isn’t whether governments are now engaged in unjustified suppression of dual-use research, it’s whether the frameworks that are being developed to guide the regulatory handling of this research guard against this danger. As we explain below, there are approaches available that *do* guard against the threat, but which are also consistent with recognising a strong case in favour of regulating dual-use research in those exceptional cases where the risks are particularly extreme.

³⁴ This is a widely-accepted approach among free speech scholars. As Brison says, while some scholars “argue that certain reasons for restricting speech are always impermissible”, even those who reject this hard-line approach think the protection of speech should get a special weighting in the policy calculus used to assess speech-restrictive policies. “Even on the balancing approach”, she explains, “the value of free speech is taken to justify weighing interests with ‘a thumb on the scales’ in favor of speech”; Brison, “The autonomy defense of free speech”: 319.

³⁵ This strategy broadly tracks the “clear and present danger” test in First Amendment jurisprudence; see e.g. Zechariah Chafee Jr., “Freedom of speech in war time”, *Harvard Law Review* 32/8 (1919): 932-73.

4. Protecting scientific autonomy in practice

First consider the practical problems. Where government regulatory agencies oversee and enforce the censorship of dual-use research, what institutional measures could we use to prevent them from abusing their powers – but without just reverting to absolutism about scientific autonomy? And how could we deal with perverse incentives on the part of the regulator to overestimate the potential harms of dual-use research and underestimate the benefits?

The first institutional design principle that we can borrow from free speech theory is a staple of the high school civics curriculum: the separation of powers. An enduring concern in free speech theory is the fact that in relation to forms of expression like protest and investigative journalism, governments are incentivised towards adopting policy that is at cross-purposes with the interests of society at large. To take an obvious example, government actors have an incentive to suppress dissent, rather than fostering the kind of open democratic debate that better serves the general social interest in transparent and accountable government. Elected legislatures are most directly subject to these bad incentives. Their effects can be counteracted, then, by vesting the power to strike down laws that unduly impinge upon free speech in a judicial branch of government. A judiciary that is independent from the legislature is less susceptible to the electoral pressures that generate the bad incentives for legislators.

This basic principle of the separation of powers could be applied in the governance of dual-use research. Insofar as there is pressure on regulatory agencies to overestimate the dangers of dual-use research, in order to avoid liability in the event of a bad outcome, states should seek to increase the degree to which their regulatory institutions are (i) independent from agencies of the state that are directly or indirectly subject to short-term electoral pressures, and (ii) staffed in a manner that minimises the relevant institutions' structural incentives to overestimate the probability-weighted costs of potentially harmful research. Creating a governance body that operated under a separation of powers would deviate from existing governance or advisory mechanisms such as the NSABB. The NSABB is a federal advisory committee whose members are selected, following the *Federal Advisory Committee Act* (FACA), by the agency that oversees the committee.³⁶ Board members are “Special Governmental Employees” that can be dismissed at the government's behest. One result of this lack of separation between the National Institutes of Health (NIH), which both coordinates the NSABB *and* is a primary sponsor of GOF research in the vein of the ferret H5N1 studies, is that the NIH can shape the membership of the board that is meant to provide advice on the regulation of research it supports. This capacity became controversial in 2014, when

³⁶ Pub. L. 92-463, 86 Stat. 770, enacted October 6, 1972.

the NIH dismissed the inaugural members of the NSABB just as the biosafety concerns about GOF research were becoming more apparent.³⁷

The second institutional design principle that we can borrow from free speech theory – in this case, borrowing from the form of free speech we find in U.S. Constitutional jurisprudence – is the use of differential review standards, incorporating the use of ‘strict scrutiny’. To elaborate: legal restrictions on free speech in American law are classified, in the process of judicial review, as being either content-based or content-neutral. Roughly, restrictions are content-*based* if the content of the expressive acts that are being restricted is appealed to as part of the justification for the restriction; they’re content-*neutral* if the content is irrelevant to the rationale behind the restriction – that is to say, if it is merely the time, place, or manner of the speech that is invoked in the rationale for its restriction. Content-based restrictions are subject to a strict scrutiny standard of review, on which they are struck down unless they are (i) a means to a compelling social interest, (ii) narrowly-tailored to the pursuit of that interest, and (iii) the least speech-restrictive means available to the government in the pursuit of that interest. Content-neutral restrictions, by contrast, are subject to the less stringent review standard of intermediate scrutiny – roughly, restrictions on speech satisfy the requirements of this intermediate standard provided that they further an important social interest.³⁸

The purpose of this two-tiered approach can be stated simply. We have very strong reasons to oppose the government’s imposition of content-based restrictions on speech. We also have strong – although significantly less strong – reasons to oppose the government’s imposition of content-neutral restrictions on speech. In neither case are the reasons unconditional or absolute in strength. What we need overall, then, is a system of checks and balances that leaves open the possibility of regulating speech in any given case, but which forces the government to satisfy a very exacting justificatory standard if it wants to enact a content-based restriction, and a non-trivial but less exacting standard if it wants to enact a content-neutral restriction. These checks and balances should apply to all areas of government policy, whether it’s enacted through legislation or through the operations of the state’s administrative and regula-

³⁷ Jon Cohen, “Updated: U.S. biosafety panel to come out of hibernation with new members” *Science Insider*, July 15, 2014. <http://www.sciencemag.org/news/2014/07/updated-us-biosafety-panel-come-out-hibernation-new-members>; retrieved March 22, 2017. The text of the dismissal letter can be found at Kelly Hills, “Remaining inaugural members of NSABB dismissed last night”, *Life as an Extreme Sport* (July 14, 2014); www.kellyhills.com/blog/remaining-inaugural-members-of-nsabb-dismissed-last-night/#sthash.zyIRVhBu.dpuf; retrieved March 22, 2017.

³⁸ For discussion see e.g. Jerome A. Barron and C. Thomas Dienes, *First Amendment Law*, 3rd Edition (St. Paul: Thomson West, 2004): 29-32; Rodney A. Smolla, *Free Speech in an Open Society* (New York: Vintage, 1993): 233-35.

tory agencies. The purpose of these checks and balances is to create robust safeguards against, for example, the suppression of dissent and unpopular opinions, while still allowing for such suppression at moments of severe, exceptional risk (e.g. in wartime), and not unduly obstructing the government in its legitimate and routine regulation of forms of speech that do not require special constitutional protections (e.g. in the regulation of false advertising).

Differential review standards can and should be applied in the governance of dual-use research. Scientific regulatory agencies should be encouraged to develop review systems that differentiate the various modes of restriction upon scientific research (e.g. by distinguishing information-suppression approaches from de-funding approaches), and require more intrusive modes of regulation to satisfy a more stringent standard of review before being enacted. Indeed, the strict *versus* intermediate review methods from U.S. constitutional jurisprudence may be a useful model for the regulatory systems we need in the dual-use context.³⁹ We could require, for example, that instances of the more intrusive modes of scientific regulation be disallowed, unless they are (i) are a means to a *compelling* social interest, e.g. the prevention of a serious pandemic illness, (ii) narrowly-tailored to the pursuit of that interest, and (iii) the least restrictive means available to the regulatory agency in the pursuit of that interest.

Most censorship of dual-use research is likely to be content-based, given that the misuse of content is what makes the research dangerous. But this distinction could still inform our approach to censorship by highlighting the role that ‘time’ and ‘place’ play in different instances of dual-use research. In 2014, for example, a novel strain of the botulinum toxin – the most lethal toxin on earth – was discovered, for which no there was no known effective antitoxin. The researchers agreed to suppress information about the sequence until such time as an antitoxin could be developed or discovered.⁴⁰ The information on the sequence was published in June 2015, with the discovery of an existing antitoxin that effectively neutralised the new strain.⁴¹ This is closer to ‘time and place’ censorship of dual-use research, compared to the

³⁹ In the context of scientific research in the United States, some of this argument has been rehearsed in discussion of laws designed to limit funding for research on human embryos, e.g. Barry P. McDonald, “Government regulation or other abridgements of scientific research: the proper scope of judicial review under the First Amendment”, *Emory Law Journal* 54/2 (2005): 979-1092; Steve Keane, “The case against blanket First Amendment protection of scientific research: articulating a more limited scope of protection”, *Stanford Law Review* 59/2 (2006): 505-50.

⁴⁰ David C. Hooper and Martin S. Hirsch, “Novel *Clostridium botulinum* toxin and dual use research of concern issues”, *The Journal of Infectious Diseases* 209/2 (2014): 167.

⁴¹ See Susan E. Maslanka, Carolina Lúquez, Janet K. Dykes, William H. Tepp, Christina L. Pier, Sabine Pellett, Brian H. Raphael, Suzanne R. Kalb, John R. Barr, Agam Rao, and Eric A. Johnson, “A novel Botulinum neurotoxin, previously reported as sero-

proposed censorship of the 2011 GOF studies, in which the motivating concern behind the NSABB's conduct was the persistent threat of a genetically modified, mammalian transmissible strain of highly pathogenic avian influenza. The general point, at any rate, is that a differential review process is an institutional measure that supports the kind of regulatory structure that we need around dual-use research. It creates a structure that permits the state's intervention in especially high-risk cases, while also maintaining a robustly-protected communicative space in which researchers can pursue their legitimate aims in the normal run of cases, without having to constantly worry about indiscriminate regulatory intrusion.

5. Protecting scientific autonomy in principle

Let's move on to the theoretical problem. As noted above there is a danger that scientific autonomy will be unduly constrained if regulators, when determining whether communication around dual-use research should be restricted, apply an overly simplistic cost-benefit framework of assessment. Instead we need a framework that can support the consequentialist reasoning used to justify the suppression of dangerous research in exceptional cases, but which at the same time offers assurance to the researcher that, in the normal run of cases, her research will be protected against indiscriminate intervention. In this section we show how the contractualist free speech framework formulated by Thomas Scanlon may be useful to this end.⁴²

The central example from Scanlon's account provides a valuable analogical reference point for discussion of dual-use research. Scanlon asks us to consider laws that ban the dissemination of instructions about how to construct explosive devices using everyday materials. On one hand, it seems intuitively plausible that a government – even a liberal government committed to pluralism, freedom of conscience, and free speech in general – is justified in imposing such bans, given the serious dangers that result from allowing this information to be freely distributed. On the other hand, as Scanlon says, free speech princi-

type H, has a hybrid-like structure with regions of similarity to the structures of serotypes A and F and is neutralized with serotype A antitoxin", *The Journal of Infectious Diseases* 213/3 (2016): 379-85; Paul Keim, "A novel botulinum neurotoxin and how it tested our scientific institutions", *The Journal of Infectious Diseases* 213/3 (2016): 332-34.

⁴² Thomas Scanlon, "A theory of freedom of expression", *Philosophy & Public Affairs* 1/2 (1972): 204-26. In later work Scanlon modified his position. The government's duty to refrain from restricting radical political dissent and the like is grounded, so he later argued, in our interest in autonomy as a good to be promoted, as opposed to autonomy as a source of categorical rights; see Thomas Scanlon, "Freedom of expression and categories of expression", *University of Pittsburgh Law Review* 40/4 (1978): 519-50. References to 'Scanlon's argument' here refer to the view advanced in the 1972 paper, and this shouldn't be understood as his final or official view.

ples cannot generally condone the suppression of any speech that creates a risk of violence. This is because many forms of political speech – especially radical and militant criticism of the government, like, say, revolutionary Marxist discourse – have the potential to foment violence. It is part of the core purpose of free speech principles, on any standard characterisation, to protect these types of political speech against government suppression. If the risk of violence *per se* can be used to justify their suppression, then free speech protections will be completely eroded.⁴³

However, Scanlon believes that we can justify the suppression of bomb-making instructions in a manner that does not force us to this unwelcome conclusion. If our bans on bomb-making instructions are strictly tailored to the aim of harm-prevention (i.e. they are not written in a way that allows them to be used as a cover for blatantly partisan political censorship), people would have no reasonable grounds to refrain from contracting into a general system of rights and liberties that allowed for such bans.⁴⁴ They do not impair anyone’s ability to be a self-governing agent in a community of equals, and they do not insulate the government from dissent or stifle public discourse about matters of fact and opinion. Although they do control people’s communication, such laws are less like the ideological censorship that free speech principles paradigmatically disallow, and more like restrictions on dangerous goods.

By contrast, Scanlon argues, people *would* have reasonable grounds to refrain from contracting into a system of rights and liberties that allowed for bans on advocating revolutionary violence, even if such bans were similarly likely to reduce the incidence of violence. There is a relevant difference between the

⁴³ One difficult case in this area is incitement to acts of terrorist violence directed against civilians. There is some reason to think that where this particular kind of political advocacy is concerned, some restriction would be permitted by properly-formulated free speech principles, notwithstanding the overtly political character of the speech. For discussion of this difficult case see Eric Barendt, “Incitement to, and glorification of, terrorism” in Ivan Hare and James Weinstein (Eds), *Extreme Speech and Democracy* (Oxford: Oxford University Press, 2010): 445-62.

⁴⁴ To conduct oneself ‘reasonably’, in the sense we have in mind, in essence, is to conduct oneself in a way that expresses a willingness to abide by fair terms of social cooperation. This characterisation of reasonableness was originally put forward in W. M. Sibley, “The rational versus the reasonable”, *Philosophical Review* 62/4 (1953): 554-60, and was famously employed in John Rawls’s theory of justice as fairness; see e.g. *Political Liberalism* (New York: Columbia University Press, 1993). There is a larger theoretical problem that’s raised by an appeal to this notion of reasonableness, which is about how claims about what is ‘reasonable’ can avoid being parochial or partisan, and whether a normative theory that assigns a major role to judgements about reasonableness can guard against this sort of parochiality; see e.g. Bruce W. Bower, “The limits of public reason”, *The Journal of Philosophy* 91/1 (1994): 5-26; Marilyn Friedman, “John Rawls and the political coercion of unreasonable people” in Victoria Davion and Clark Wolf (Eds), *The Idea of a Political Liberalism: Essays on John Rawls* (Oxford: Rowman and Littlefield, 2000): 16-33. There is a legitimate theoretical challenge to be addressed here. As we explain below, though, the judgements about non-reasonableness that are involved in our application of the contractualist framework, in this context, are ones that are accepted by all participants in the dual-use debate. This means that in the context of this debate, worries about parochiality in our judgements about what counts as reasonable have no real purchase.

cases, because the second kind of case impinges on people’s entitlement to decide for themselves what to believe, how to respond to political advocacy, and what ideas and exhortations are fit to be treated seriously. As Ronald Dworkin puts it, in a discussion that echoes this element of Scanlon’s account, a government “insults its citizens, and denies their moral responsibility, when it decrees that they cannot be trusted to hear opinions that might persuade them to dangerous or offensive convictions”.⁴⁵

For Scanlon, the discrepancy between these two kinds of cases lays the foundations for a contractualist free speech principle. Dangerous speech cannot be legitimately restricted if the underlying rationale is one that a self-governing people can reasonably refuse to contract into, on grounds like those outlined above.⁴⁶ The government that censors instructions that help increase the number of bombs being built in people’s sheds is different, on his account, to the one that censors speech that may persuade people to join violent political movements.

Scanlon’s account indicates how, when it comes to regulating dual-use research, we can have a framework that supports the consequentialist reasoning used to justify the suppression of dangerous research in extreme cases, while at the same time reassuring the researcher that her work will not be subject to regulatory interference in the normal run of cases. There are all sorts of tenuous bases on which academic research programs could be deemed dangerous. This is most evident outside the realm of the natural sciences. For instance, research in the humanities and social sciences can help to foster social movements that destabilise the social and political status quo.⁴⁷ But even in the natural sciences ideologically dubious claims about the dangers of research are also a genuine problem that needs to be guarded against. For instance, certain political actors may have spurious or ideologically-grounded reasons for regarding some research in the natural sciences as dangerous or politically inconvenient. For example, in 2003 officials at the U.S. Environmental Protection Agency tried to suppress data pertaining to global historical

⁴⁵ Ronald Dworkin, “Why must speech be free?” in *Freedom’s Law: The Moral Reading of the American Constitution* (Cambridge Massachusetts: Harvard University Press, 1996): 200.

⁴⁶ In a representative passage, Scanlon says his principle of free speech, “as a general principle about how governmental restrictions on the liberty of citizens may be justified, is a consequence of the view, coming down to us from Kant and others, that a legitimate government is one whose authority citizens can recognize while still regarding themselves as equal, autonomous, rational agents”. And thus, he says, although “it is not a principle about legal responsibility... [it] has its origins in a certain view of human agency from which many of our ideas about responsibility also derive”; Scanlon, “A theory of freedom of expression”: 214.

⁴⁷ For instance, a number of critical social movements – feminism, queer activism, transgender activism, and radical environmentalism, to name a few – have been nurtured and developed to a significant degree by the work of academic researchers in humanities and social science disciplines. (This isn’t to suggest that academic research was the primary driver of these movements, just that it was a non-negligible influence.)

temperatures, human causes of global climate change, and the harmful consequences of such climate change, in an attempt to undermine claims that supported government action to mitigate the causes of anthropogenic climate change.⁴⁸ According to the logic of Scanlon's contractualist free speech framework, any justifications of this kind for restricting 'dangerous' research would be straightforwardly disallowed.

By contrast, there are some serious risks that all reasonable parties recognise *as* serious risks, like widespread access to bomb-making instructions, or access to technical data that could be used to synthesise bioagents and trigger a global pandemic. Restrictions on communication aimed at minimising these risks are permitted, in principle, by a contractualist free speech framework like Scanlon's. In cases of radical political advocacy, people can only agree to let the government protect them from this 'dangerous speech' by relinquishing a sense of themselves as self-governing agents. To ask the government to protect us from such speech is to ask the government to treat us as children. But people need not relinquish any sense of themselves as self-governing, free and equal agents in order to accept restrictions on access to information whose dangerousness owes to the fact that it can help to facilitate intentions to carry out a program of indiscriminate violence and killing. These sorts of grave risks are different in kind from outcomes that can only be characterised as 'grave risks' if one is judging things from a specific politically- or ideologically-slanted standpoint (e.g. things like the disorder of a revolutionary political uprising, or the political fallout from acknowledging the existence of climate change). Restrictions on communication aimed at minimising *these* sorts of 'dangers' would not be permitted by a contractualist free speech framework like Scanlon's.

It is true that the application of a framework like Scanlon's requires us to draw distinctions between which kinds of political views qualify as reasonable and which ones do not, and that such distinctions embed normative judgements or ideals that are contestable in principle.⁴⁹ In particular, when this framework is used as a rationale for censoring things like dual-use research or bomb-making instructions, it builds in the assumption that all political views that accept indiscriminate, politically-motivated killing, directed against the state or other citizens, can be relegated to the domain of the unreasonable and sanctioned accordingly. This kind of assumption could be argued against, in principle. Having said that, the entire debate in this arena is premised on an equivalent assumption. It is premised on the view that at least some high-risk dual-use research is too dangerous to be freely conducted and/or circulated, because of its potential application in the development of bioagents, and that complaints from those who regard

⁴⁸ Tania Simoncelli and Jay Stanley, *Science Under Siege: The Bush Administration's Assault on Academic Freedom and Scientific Inquiry* (New York: American Civil Liberties Union, 2005).

⁴⁹ See note 46.

this outcome as disadvantageous to their political agenda are grounded in an unreasonable politics, and can be disregarded accordingly. Someone who takes issue with this assumption is not merely objecting to a contractualist theory of free speech. Rather, they are objecting to the whole underlying rationale for regulating dual-use research. As explained above, we think there is a decisive case in favour of such regulation, at least in extreme instances.

The point of a contractualist framework is not to embed a controversial notion of reasonableness in this debate. This point is to accommodate the judgement that some dual-use research may be too dangerous in its potential applications, and hence must be censored, but *without* reverting to a simplified cost-benefit framework, which would too easily allow concerns about other kinds of ‘dangers’ to be cited to justify the censorship of other kinds of speech. The contractualist framework sets principled standards – based on what kind of system of rights and liberties people could reasonably refrain from contracting into – that limit which kinds of outcomes a regulatory agency can cite as ‘dangers’ in order to justify gagging, de-funding, or otherwise coercively interfering with a dual-use research program. In comparison to a simple cost-benefit framework of assessment, the appeal to these standards makes it harder for regulatory agencies to use claims about ‘the danger of research’ as a way of smuggling in blatantly partisan judgements about which kinds of research are of social merit more generally. Appeal to these standards will also reduce the degree of uncertainty about which lines of research will be deemed unacceptably dangerous by a regulatory agency, thereby helping to mitigate the chilling effects created by this uncertainty. These are the main advantages of adopting a contractualist framework for thinking about the limits of justifiable censorship in this area.

In a number of ways, then, a contractualist framework for adjudicating decisions about the regulation of dual-use research would – together with appropriate institutional measures, like those discussed above – support the autonomy of scientific researchers in this arena. Crucially, though, they would support their autonomy in a way that does not just collapse back into an implausible form of absolutism about the rights of scientific researchers. Under the contractualist framework, censorship of genuinely dangerous dual-use research could be entirely justifiable and appropriate, in-principle. The point is that the rationale for such regulation cannot be about micro-managing or pre-emptively guiding the outcomes of research, or supporting the government’s favoured politics. Regulatory control of communication around the most dangerous forms of dual-use research does not betoken an acceptance of regulatory agencies having a more extensive role as arbiters of good or worthwhile scientific practice.

6. Regulating dual-use research: pragmatic concerns

We have shown how resources from free speech theory can help to reconcile respect for scientific autonomy with coercive government regulation of dual-use research. But we have not yet said anything to address doubts about the effectiveness and consequences of coercive government regulation in this arena. In this closing section we address three types of doubts. The first relates to the efficacy of censorship: given that scientific research is easy to publish online, with or without the support of a journal, it may be nearly impossible to effectively censor a researcher who is determined to publish their work. The second concern is that in censoring research, scientists might risk being “scooped” by their competitors who are not subject to the same rules, e.g. those working in jurisdictions with less stringent regulatory policies. The third pragmatic concern about censorship is that, given the ubiquity of life sciences research and innovation worldwide – much of it conducted by individuals and private firms – there may be no stage in the ongoing research process at which government review for the purposes of censorship could adequately capture the research which we are worried about.⁵⁰

A common theme underlying these concerns is about the efficacy of censorship when applied to the modern life sciences. We acknowledge that the information age makes censorship – even if permissible, all things considered – difficult. Feasibility should absolutely act as a constraint against censorship in the same way as it does in other liberty-impairing acts.⁵¹ But feasibility comes in degrees, and the degree of success we require to justify censorship of dual-use research must be weighed against the costs of inaction. If some piece of dual-use research posed a global catastrophic risk that would permanently harm humanity in an irreversible way (e.g. by killing billions of people), censoring that research would be justifiable in principle, even if our methods only provided limited means of containing the information (and thus preventing disaster).

Moreover, some regulatory policies would satisfy a reasonable definition of feasibility for a range of dual-use research, and would thus be permissible assuming they satisfied our other conditions. For example, the 2012 and 2017 U.S. Government policies on dual-use research both recommend that, if deemed necessary by agencies, dual-use research should be moved to funding agencies that had the power to classify scientific research (e.g. from the Department of Health and Human Services, to the Department of

⁵⁰ We thank two anonymous reviewers from this journal for raising these additional objections.

⁵¹ See e.g. James F. Childress, Ruth R. Faden, Ruth D. Gaare, Lawrence O. Gostin, Jeffrey Kahn, Richard J. Bonnie, Nancy E. Kass, Anna C. Mastroianni, Jonathan D. Moreno, and Phillip Nieburg, “Public health ethics: mapping the terrain, *The Journal of Law, Medicine & Ethics* 30/2 (2002): 170-78.

Defence). Classified research is performed around the world and, at least in those cases that have been subject to study, “classified communities” can permit free scientific inquiry *qua* the selection of interesting questions and communication with peers to a degree that is tolerable to many scientists.⁵² While some information risks remain, from malevolent actors inside or outside the community (e.g. the risk of a cyberattack), the historical record seems to indicate that classified science is not only effective, but can be very productive. This suggests that there are applicable methods of information containment that would provide feasible means of regulating dual-use research.

These methods may have their own costs, of course. This dovetails into the second concern: that censorship may disincentivise the conduct of science. Making research classified definitely alters the incentive structure for researchers, including priority over discoveries, the potential to secure intellectual property rights or patents, or engage in entrepreneurial activity. In the physical sciences, for example, the classification of nuclear secrets has been charged with undermining of civilian nuclear power efforts, by restricting information to only those working on national security issues, or firms large enough to afford the logistical burden of attaining clearances for the purpose of developing civilian nuclear power.⁵³

However, our proposals would not involve restricting an entire discipline, as opposed to only a fairly limited class of research programs in the life sciences. In the case of GOF, it may be possible to not only restrict information, but also to limit the conduct of this research to a smaller number of locations. In testimony before the U.S. House Energy and Commerce Committee in 2014, Richard H. Ebright recommended that the number of high-containment, Biosafety Level 3 and 4 laboratories in the U.S. be reduced to 25-50, from the existing level of approximately 1,500.⁵⁴ We could envisage an environment in

⁵² Peter J. Westwick, “Secret science: a classified community in the national laboratories”, *Minerva*, 38/4 (2000): 363-91.

⁵³ Robin Cowan, “Nuclear power reactors: a study in technological lock-in”, *The Journal of Economic History* 50/3 (1990): 541-67. Cowan’s studies of the lock-in of nuclear power – lock-in that continues today – don’t solely focus on the role of classification. He also blames regulatory capture of the industry following rare but high-profile disasters such as the meltdown of the Three Mile Island reactor and its subsequent investigation. This is a significant addition, given the current debate on gain of function research is also focussed on rare but potentially calamitous accidents. Other work on the evolution of physics research, however, has noted that there were broad effects on physics research during the Cold War that altered research trajectories through their use of classification as a means of containing information and dictating the kinds of researchers that had access to the community. See e.g. Helge Kragh, *Quantum Generations: A History of Physics in the Twentieth Century* (Princeton: Princeton University Press, 1999); Paul Forman, “Behind quantum electronics: national security as basis for physics research in the United States, 1940-1960”, *Historical Studies in the Physical and Biological Sciences* 18/1 (1987): 149-229.

⁵⁴ Written Testimony of Richard H. Ebright, submitted for the record to the House Energy and Commerce Committee Subcommittee on Oversight and Investigations, for the hearing “Review of CDC Anthrax Incident” July 16, 2014; see <http://docs.house.gov/meetings/IF/IF02/20140716/102479/HHRG-113-IF02-Wstate-EbrightR-20140716.pdf>, retrieved 14 September, 2017.

which the most dangerous GOF studies were confined to a small number of laboratories---or, perhaps, even just one high security facility.

Incentives can be replaced in these environments. For a start, the opportunity to work on very select problems can be its own reward. Classified laboratories create incentives for scientists by offering employment or salary benefits, access to research resources, or opportunity to work on problems that cannot be worked on anywhere else. We concede that in some cases scientists may give up some important parts of the scientific process, such as securing priority or public recognition of work, while being compensated in other ways.

The final objection raises questions about how and when we should regulate research of concern. With the increasingly dominant role of private research funding, and the decreasing economic and technical barriers to access the life sciences, it is likely that the private sector will eventually be (if it is not already) a significant driver of dual-use research.⁵⁵ This pertains to broader concerns around the regulation of scientific research. The U.S. Common Rule, as a federal rule that only applies to research or institutions in receipt of federal funding, has encountered similar issues.⁵⁶ Appropriate regulation of private research, in general, is an important unresolved issue warranting further analysis.

Any regulation is likely to miss something. A better framing of the objection asks what a sufficiently robust (and beneficial) policy that justifies the costs of censorship would look like. Answering this is largely an empirical problem, and one that plagues the larger policy debate on dual-use research. Existing U.S. regulations that address dual-use, in some cases, only do so in the context of pathogens listed under the Select Agents Program, which describes biological agents and toxins that pose greatest risk of use by terrorists.⁵⁷ Yet avian influenza, while being a Select Agent for purposes of the U.S. Department of Agriculture, is *not* a select agent for the Department of Health and Human Services. Nor is the horsepox virus, which was the subject of a controversial dual-use experiment in 2017, when scientists showed how to synthesise the virus from base sequences. This experiment was criticised for providing a roadmap to

⁵⁵ On the changing role of federal and private research funding in the US, see Homer A. Neal, Tobin L. Smith, and Jennifer B. McCormick, *Beyond Sputnik: U.S. Science Policy in the Twenty-First Century* (Ann Arbor: University of Michigan Press, 2008). On the change in economic and technical barriers to accessing the life sciences, see Nicholas G. Evans and Michael J. Selgelid, "Biosecurity and open-source biology: the promise and peril of distributed synthetic biological technologies", *Science and Engineering Ethics* 21/4 (2015): 1065-83.

⁵⁶ David Hunter and Nicholas G. Evans, "Facebook emotional contagion experiment controversy", *Research Ethics* 12/1 (2016): 2-3.

⁵⁷ Office of Science Policy, *United States Government Policy for Institutional Oversight of Life Sciences Dual Use Research of Concern* (Bethesda: National Institutes of Health, 2014).

synthesise smallpox, and as a clear instance of problematic dual-use research.⁵⁸ In 2010, the NSABB and others noted that pathogen-based regulation of dual-use research is inappropriate in the emerging life sciences, and that functional analysis (i.e. what properties do specific sequences confer on organisms) would be preferable.⁵⁹ To date, this has not been implemented at a national level, although life sciences communities have begun to investigate such a “post-taxonomic” risk analysis as part of their internal regulations.⁶⁰

In summary, the position that we have defended is not premised on the implausibly optimistic notion that we might be able to fully protect societies against the risks associated with dual-use research via implementation of scientific censorship (in exceptionally risky cases). It is premised on the much more modest notion that when it comes to managing relevant risks, some government regulatory controls (including limits on dissemination of information) are better than nothing. The value of scientific autonomy, as far as it goes, is insufficient to outweigh the case for imposing regulatory controls on dual-use research that carries catastrophic risks. When considering how these controls would operate, it is entirely appropriate to be concerned about a rationale for government censorship that has based on claims about the danger that the speech in question poses. The abuse of government power is a genuine and legitimate danger. But there are institutional design principles and normative frameworks from free speech theory that can be used to augment a risk-based rationale for regulating dual-use research, in a way that makes our framework less susceptible to such abuse. The fact that regulatory oversight of dual-use research will not necessarily be as efficacious as we would ideally like is no reason to refrain from doing what we can to mitigate danger. The risks that arise from at least some cases of dual-use research may be so great that even limited reduction would be worthwhile and justifiable.

⁵⁸ Kai Kupferschmidt, “How Canadian researchers reconstituted an extinct poxvirus for \$100,000 using mail-order DNA”, *Science* (July 6, 2017).

⁵⁹ National Science Advisory Board for Biosecurity, *Addressing Biosecurity Concerns Related To Synthetic Biology* (Bethesda: National Institutes of Health, 2010).

⁶⁰ Piers Millet, “Moving past taxonomy-based screening: matching measures to risks”, *Bio.secure* (April 1, 2015); http://biosecu.re/biosecure/writing/Entries/2015/4/1_Moving_past_taxonomy-based_screening__matching_measures_to_risks.html, retrieved September 20, 2017. One of the authors is a member of the Safety Committee of the International Genetically Engineered Machine Competition, which implements Millet’s screening method. Little public documentation currently exists on this method, outside of Millet’s sporadic writings.