Global Public Policy and Governance (2022) 2:266–282
https://doi.org/10.1007/s43508-022-00045-y

RESEARCH ARTICLE

Governing emerging technologies—looking forward with horizon scanning and looking back with technology audits

Henry T. Greely[1]

Received: 17 July 2022 / Accepted: 16 August 2022 / Published online: 29 August 2022

Abstract
Many technologies, particularly in the life sciences, are subject to governance regulation. Discussion of such governance tends to focus on limits on research and experimental uses and then on a decision whether, and how, to allow wide use. Two other important aspects of technology governance, one earlier and one later, have been discussed less: anticipation of technologies with important effects and the monitoring of the actual effects of adopted technologies. This article will analyzes those points of the governance process and proposes a plan to improve their functioning. The article argues first that the world would benefit from a more visible and influential approach to spotting and analyzing emerging technologies through a high profile "Horizon Scanning Group." It then proposes a more formal approach to monitoring and assessing new developments through "Technology Audit Groups." The article's third section discusses complicated organizational issues surrounding these proposed groups and proposes some specific approaches.

Keywords Technology governance · Horizon scanning · Forecasting · Monitoring · Technology auditing

1 Introduction

In Kurt Vonnegut's novel, SLAUGHTERHOUSE-FIVE, the human protagonist learns the ways of the alien Trafalmadorians, to whom the whole past, present, and future is fully visible—indeed, is in immediate existence (Vonnegut, 1969). Absent a Trafalmadorian ability to know the future, there can be no perfect answer to technology governance, and the problem is worse than just ignorance. David Collingridge set out a basic dilemma in 1980: At the stage when steering a technology is easy, we

✉ Henry T. Greely
hgreely@stanford.edu

1 Center for Law and the Biosciences, Stanford University, Stanford, CA, USA

do not know how it needs to be controlled; when we have seen it develop enough to see its problems, entrenched actors and expectations make change very difficult (Collingridge, 1980).

I am not a scholar of technology assessment and governance in general; I study the biosciences and the ethical, legal, and social challenges they present. The biosciences are just one kind of technology and although important, they may not represent well how other technologies develop. However, I hope some insights from that area of science may be helpful, in bioscience policy and beyond.

From the perspective of the biosciences, the life cycle of a policy responses to new technologies can be grouped into four stages: the initial recognition of a possible new technology that might have policy implications, the assessment of preliminary research with the technology, decisions about adoption and regulation of the technology, and monitoring of the actual effects of the technology. In the biosciences, the second and third areas are usually well-covered.

Pre-adoption research with humans will be reviewed by ethics committees with different names in different countries ("institutional review boards" or IRBs in the United States), research with non-human animals will usually be reviewed by animal care and use committees, and yet other committees will try to ensure that the research is done in ways that will be safe for those working inside and outside the research settings. Plus, the research will be contingent on funding and on decisions by the institutions, where the research will take place, whether commercial, academic, or governmental. Each will apply its own set of risk/benefit calculations to a decision to go forward.

Then, in the biosciences, unlike many areas of the economy, new technologies usually require some kind of official approval before they can be marketed or widely adopted. Health-related biotechnologies require approval by various drug, biological product, and medical device bodies in most countries. Other biotechnologies will require environmental or food safety approvals.

But we pay too little attention to the first and fourth steps: anticipation and monitoring. This article is fundamentally a proposal, or perhaps a thought experiment, for organizations and policies that would improve technology governance by bolstering those two steps. First, it argues that the world would benefit from a more formal approach to spotting and analyzing emerging technologies through a high profile "Horizon Scanning Group." It then proposes a more formal approach to monitoring and assessing new developments through "Technology Audit Groups." The article's third section will discuss complicated organizational issues surrounding these proposed groups and proposes some specific approaches. Although it is drawn largely from observations of developments in the biosciences, I believe it may be useful in other fields.

2 Scanning the horizon

This section discusses efforts to spot emerging technologies early. The term "horizon scanning" is often applied to these efforts and is said to have appeared first in the title of a 1995 business publication by Microsoft founder Bill Gates, although presumably it has deeper maritime roots as people, on the shore or on ships, would scan the horizon, looking for other ships or land (Gates, 1995).

What I mean by horizon scanning is an effort to spot plausible but not immediate technological possibilities before they are adopted and then to assess which ones, if any, are likely to become important. This, in a way, is like predicting Atlantic hurricanes, where low pressure areas off the west coast of Africa are identified and watched well before they become either devastating Category 5 storms or quickly passing rain squalls—or, most often, something in between. Like early hurricane predictions, they would often be wrong; plausible technologies, like low pressure systems, will not all develop to be world changing—or develop at all.

People scan technical horizons for many reasons, including notably to look for the commercial promise of possible technologies, but this article is concerned with their more public consequences—their ethical, legal, social, and policy implications. The hope is that with advance attention, these technologies might be adopted in ways that minimize their risks and maximize their benefits (although this is more a hope than an empirically demonstrable proposition).

The very early stage of the examined technologies distinguishes horizon scanning from some other work before the adoption of particular technologies. For example, it is not like an environmental impact statement, required by federal or state laws in the United States before undertaking actions that might affect the environment. Those try to map out in detail the likely environmental effects of a very specifically defined project, such as a building a two lane asphalt public road along a defined route. Neither would it be like review by an Institutional Review Board or its equivalent, assessing the ethical propriety of a research with human subjects. An IRB looks at a specific protocol that says exactly what will be done to or with the human research participants and after what kinds of consent. Horizon scanning efforts should set out why and how a technology might have eventually have significant effects on society, but not to specify steps in its adoption.

Perhaps most importantly, to be useful, horizon scanning efforts must be used. Their results must be seen, discussed, and, as appropriate, followed up. Otherwise, they are only footnotes in someone's future historical research.

This section will begin with a recent example of an ad hoc but successful horizon scanning effort in human germline genome editing. It will then look briefly at the wide range of existing horizon scanning endeavors and methodologies. It will finally discuss some shortcomings in current horizon scanning approaches.

2.1 A largely successful example of horizon scanning

In many areas new technologies are carefully examined and weighed for their ethical, legal, and social implications, with much of the investment coming from foundations, government grants, or from work by scholars who do not need outside funding. Work on the ethical implications of in vitro fertilization, of cloning, of human embryonic stem cell research, of functional magnetic resonance imaging in humans, and of human germline genome editing did happen, but, in most cases, only after announcements led to a shock: from Louise Brown and Dolly the sheep to James

Thomson's announcement of his creation of human embryonic stem cells to numerous neuroscience studies that garnered headlines by announcing the discovery of the brain location for true love, the feeling of mystical communication with God, or aggressive behavior. Human germline genome editing is a useful exception—and, to some extent, demonstrates what I think is a good horizon scanning endeavor.

"Genome editing" is the process of making intended changes in an individual organism's DNA sequence, its "genome."[1] Human "germline" genome editing involves changes that do not just involve the person who is edited but, through changes made to that person's eggs or sperm (their "germline"), may affect that person's descendants, potentially for millennia. People have been talking about, and worrying about, humans changing their descendants' genes for decades,[2] with serious ethical discussion beginning no later than the 1960 s, starting with religious scholars and philosophers. That discussion continued through the various stages of the "genetics revolution," with no consensus reached and no strong policies being adopted—probably because they were not needed. No one knew *how* to edit the human germline genome so the discussions remained hypothetical.

The most plausible way to edit a human's germline would be to edit the embryo that will become that human, changing the DNA of the embryo at an early enough stage that all of its cells, including the cells that will become their eggs or sperm, would be changed.[3] Living ex vivo (outside the body) human embryos became available through the development of in vitro fertilization, after the first human IVF birth in 1978. Researchers' understanding of how to maintain and manipulate such human embryos expanded over time, particularly during and after the 1990 s. There was still no good way to edit DNA, of those embryos or of anything else, but discussions of the ethical and social issues in human germline editing continued, although without urgency or resolution.

Editing DNA became more possible in the decade of the 2000 s with the development of two new gene editing tools, TALENs (which stands for "transcription activator-like effector nucleases") and Zinc Finger Nucleases. These techniques, however, are difficult, time-consuming, expensive, and uncertain. Then, in June 2012, Jennifer Doudna and Emmanuelle Charpentier published their discovery of how to use CRISPR ("clustered regularly interspaced short palindromic repeats") to edit DNA, which quickly led to the recognition that this was going to be much faster, cheaper, and easier than it had ever been before (Jinek 2012).

But it was about 18 months before appearance of CRISPR began to affect the human germline genome editing debate. In 2014 Jennifer Doudna became very

[1] The discussion in this section is taken largely from *CRISPR People* (Greely, 2021).

[2] Interestingly, Aldous Huxley's Brave New World, which is often viewed as the originator of this story, did not actually involve any gene editing. Written twenty years before the structure of DNA was discovered, the book does manipulate the abilities of people from a very early stage but through selecting the genetic parents of particular babies, giving the fetuses different treatment in their artificial wombs ("bottles"), and providing them with vastly different training after birth (Huxley, 1932).

[3] One might also approach human germline editing by editing the eggs and the sperm of a living person and using those modified germ cells to make a new embryo. This is technically more difficult than editing an embryo and has received little attention so far.

concerned about the possible uses of the tool she had helped create, spurred, in her telling, by a vivid dream. That fall, she and some of her colleagues organized a workshop to be held in late January 2015 in Napa Valley, California. By the time the meeting was held, rumors that proved to be accurate were circulating that Chinese scientists were seeking to publish research that had used CRISPR to edit human embryos, but embryos that were known not to be viable (Liang 2015). (I do not know whether those rumors existed when the meeting was being planned.) Fourteen academics attended, twelve scientists and two law professors who had worked on genetic ethics issues. (I was fortunate to be one of them.) That meeting reached a consensus surprisingly quickly—"yes" to CRISPR for research on human cells and human gene therapy aimed at somatic cells (not eggs or sperm), but "no" for human germline editing unless or until it was shown to be safe and there was societal approval. The group, with a few additional authors, published its conclusions in Science in March 2015.

The Science paper, and particularly its first author (listed alphabetically but not unfairly), David Baltimore, then convinced the U.S. National Academy of Sciences to create a project on Human Genome Editing. That project in turn led to a World Human Genome Editing Summit in Washington, D.C. in December 2015 sponsored by the U.S. National Academy of Sciences, the U.S. then Institute of Medicine (since renamed the National Academy of Medicine), the Chinese Academy of Sciences, and the U.K. Royal Society.

The U.S. bodies then published a long report on the topic in February 2017. The Nuffield Council in the U.K. produced a preliminary report on the issues in September 2016 and a longer report in June 2018. All were broadly consistent with the approach taken in the Science paper: yes for ex vivo research and for careful use in somatic cell gene therapy but no, at least for now, for use in germline editing.

But sometime in March 2018, before the U.K. report was published, Chinese scientist He Jiankui transferred genetically edited embryos into the uteruses of women who wanted to have children from them. Two of those embryos, transferred in about March 2018, gave rise to pregnancies and, in October 2018, baby girls. Their births were announced by He Jiankui just before the Second International Human Genome Editing Summit, sponsored by the two U.S. academies, the Royal Society, and the Hong Kong Academy of Sciences. The result was an uproar and, ultimately, He's criminal conviction by a Chinese court and a sentence of a large fine and 3 years in prison.

I view these events as an example of at least a partial success for an effort at something like horizon scanning. Arguably, CRISPR was already too established by January 2015 to truly be glimpsed at a horizon but it was only about then that the technique's power began to become clear. The January workshop, and the resulting March publication, went beyond saying "this is a technology to watch" to provide some interim recommendations—recommendations that have largely been followed in subsequent reports—but its most important role was to force attention to the immediacy of the issue. It helped that of its 14 participants, two, David Baltimore and Paul Berg, had not only Nobel Prizes but were among the five organizers of the famous Asilomar meeting on scientific self-regulation of recombinant DNA, held almost exactly 40 years before the Napa meeting. The issue then got picked up, publicized, discussed, and analyzed by other groups.

Those prescriptions did not deter He; he even claimed, falsely, that his research met the conditions set by the National Academies' report. However, I would argue that they provided a common background understanding that led to an almost unanimous condemnation of He and his work. If the He experiment had been revealed without any of the nearly 4 years of discussions and deliberations that started with the Napa meeting, it is unclear what the reactions to it would have been. Surely there would have been hurried efforts to discuss and analyze human germline genome editing, but perhaps, as was the case the announcement of the birth of Dolly the sheep, with less confidence…and more near hysteria.

2.2 Existing scanning the horizon: definitions, approaches, and use

Whether called "horizon scanning," "technology assessment," "foresight," or something else, much has been written about this approach—and much of it has been done. This subsection makes a very brief survey of the field and then highlights a weakness in the existing efforts.

The idea of horizon scanning is not new; after the first use by Bill Gates, Wikipedia reports it quickly became fairly widely used, first in Europe and first for health technology trends (Wikipedia, 2022a, 2022b). In addition, indeed, its use, in some ways, predates the term with earlier approaches often called "technology assessment" (Banta, 2009). Banta finds that approach in the U.S. in the 1960 s, particularly with regard to supersonic transports, environmental pollution, and genetic engineering, and he dates the first use of the term to 1967 and the U.S. House of Representative's Committee on Science and Astronautics (U.S. Congress 1967).

A few years later, the U.S. Congress created the Office of Technology Assessment ("OTA") through the Technology Assessment Act of 1972, which stated in its Sect. 3(c) that.

The basic function of the Office shall be to provide early indications of the probable beneficial and adverse impacts of the applications of technology and to develop other coordinate information which may assist the Congress.

The Office was governed by a 12-person board, half of them members of the House of Representatives and half of them Senators, and half from each of the two main American political parties. Before it was abolished by having its appropriations eliminated in 1995, as a result of control of the House of Representatives passing to a conservative Republic majority, headed by Newt Gingrich, it had produced over 750 reports on a wide range of topics (Wikipedia 2022a, 2022b).

About the time the OTA was abolished, the term "horizon scanning" came into use (presumably coincidentally). However, once it began to be used, its definition became a question. A 2015 report prepared for the European Commission defined it as the systematic outlook to detect early signs of potentially important developments. These can be weak (or early) signals, trends, wild cards or other developments, persistent problems, risks and threats, including matters at the margins of current thinking that challenge past assumptions. Horizon Scanning can be completely explorative and open or a limited search for information in a specific field based on the objectives of the respective projects or tasks. It seeks to determine what is constant what may change, and what is constantly changing in the time horizon under analysis. A set of criteria is used in the searching and/or filtering process. The time horizon can be short-, medium-, or long-term (Fraunhöfer 2015).

This broad definition has not been universally accepted. Safeguarding the Bioeconomy, a 2020 report from the U.S. National Academies of Sciences, Engineering, and Medicine notes, in a very useful chapter on horizon scanning,

The terminology around relevant tools, techniques, and processes involved in horizon scanning has yet to be standardized, which can lead to confusion. In some cases, for example, the overall process of structured reflection on the future is referred to as "horizon scanning" (UK Government Office for Science, 2013), while in others, it is termed "foresight" or "future(s) thinking" (FAO, 2013). In this report, the committee has adopted a definition similar to that used by the Organisation for Economic Co-operation and Development (OECD): horizon scanning is "a technique for detecting early signs of potentially important developments through a systematic examination of potential threats and opportunities, with emphasis on new technology and its effects on the issue at hand" (OECD, n.d.a) (NASEM 2020).

And at least some scholars have argued that horizon scanning and foresight are different things. Cuhls, for example, urges that "Foresight and Horizon Scanning are often regarded as identical concepts with identical models and purposes. However, they can be completely different activities, both in and for any kind of prospective activities as well as for strategy purposes." She sees horizon scanning as a necessary but limited first step in a broader foresight activity (Cuhls 2019).

Whatever the definition, the actual operation of what is called horizon scanning can take many different paths. One important variation involves its time frame of a horizon scanning enterprise. The report to the European Commission noted above set out three different temporal approaches to horizon scanning activities: "continuous scanning activities to keep the overview (often with regular newsletters), regular but discontinuous activities (e.g., every 5 years) and ad-hoc Horizon Scanning for a specific purpose, on demand or at a specific occasion." (Fraunhöfer 2015). OTA was an example of a continuous activity, while the various bioethics reports about human genome editing discussed earlier were largely ad hoc efforts on that particular topic, although often performed by organizations, like the U.S. National Academies or the Nuffield Council, that were in continuous operation. The Nuffield Council, for example, has recently produced annual reports of "What's on the Horizon?" (Nuffield Council, 2022).

Just deciding when or how often to do horizon scanning, and on what topic or topics, is only the beginning of ones choices. Many different methodologies exist. A 2019 systematic literature review of methodologies found, examined, and summarized about 100 relevant articles (Hines 2019). It stated that horizon scanning "generally follows a process of signal detection, filtration, prioritisation, assessment, and dissemination," each of which it then describes. The article sees artificial intelligence as possibly resolving problems in the filtration, prioritization, and assessment steps.

Another useful discussion is found in a detailed description of the horizon scanning process used in the report of the European Union's Scanning for Science and Technology Issues project (Amanatidou 2012). This paper contains a good discussion of different tools for getting input into different phases of the scanning process, including focused expert reviews, wikis, twitter, surveys, conferences, and text-mining. It classified different approaches into "semi-automated, participatory," "semi-automated non-participatory," and "manual combined."

The Fraunhöfer report to the European Commission took yet another approach to a taxonomy of horizon scanning activities, grouping them broadly into those that lead to broad reports that take time and others to dense ones that are quick. The terms it associates with the broad approaches include "participative", "overview", "for topic legitimization", "open proposals or warnings", and "continuous or part of a full process". With the dense approaches, it associates the terms "automated", "pre-defined scan field", "independent advice on topics", "on demand", and "stand-alone"..

Finally, in whatever ways it is done, horizon scanning *is* being widely used. In governments, has been explicitly tied to general policy making processes in the United Kingdom, Switzerland, the Netherlands, and Singapore (which is focusing on automating the process) (NASEM 2020). It is used in particular cases by the European Union, the United States, Germany, Sweden, Russia, and many other places, as well as by non-profit and for profit organizations that provide it for clients.

2.3 A problem with many horizon scanning efforts

The biggest problem with horizon scanning is not that it is being done too rarely, but that it seems to have little broad impact. Governments, government departments, businesses and others do various forms of horizon scanning. They may or may not act on its results, but those results seem rarely, if ever, to become part of a broader conversation—among the researchers, entrepreneurs, academics, and, most importantly, the publics who may be affected by the cloud on the horizon. If they get no attention, the best projections have little worth. The ad hoc horizon scanning of the Doudna workshop had effects on the world, because it had some highly respected members and was able to publish in a high profile setting, getting the attention of the U.S. National Academies of Sciences, Engineering, and Medicine.

What is needed is a mechanism that can, with "sufficient" accuracy, suggest a set of technologies that need to be watched because of the combination of their plausibility and the significance of their possible social and ethical effects. Such a mechanism needs some continuity of mission and membership over time as well as substantial credibility with its audiences. It also needs to be small enough to be both nimble and, perhaps more importantly, financially plausible. In addition, most important, its conclusions need to inform a public community. For that, it needs to be from a recognized authority, not a set of secret, confidential, or merely unnoticed reports from scores of groups around the world. In addition, we need the results to be widely disseminated.

What follows is a proposal, or perhaps a thought experiment, for providing such a professional and high profile source of horizon scanning. An organization that I will call a Horizon Scanning Group should be created with an initial guarantee of 10 years of funding, renewable in 5 year increments if appropriate. I have no views on the appropriate methods or mix of methods the group should use for its horizon scanning; I would leave the choice among the many options to those experts. Its mission would be to produce an annual report on a few technologies—perhaps five—it believes may need further attention and scrutiny because of a combination of their likelihood and their possible social or ethical effects.

The Horizon Scanning Group should have a minimum of ten full time professionals on its staff with expertise, mainly at the Ph.D. level, along with an appropriate support staff. The professionals would be from a wide range of fields, including biology, physics, chemistry, computer science, and engineering, along with at least one political science or legal expert. They could have a physical office or be (largely) distributed and work remotely. These full time professionals would be supplemented by working advisory boards of well-known and highly respected researchers, relevant industrialists, non-governmental organization experts, and government officials. They would also be assisted by specialists in science communication.

These experts would have several year renewable contracts, with a short probationary period after which they could only be dismissed during the term of their contracts for good cause. The professionals would spend their time reading in the relevant literatures and having discussions, small or large, with other experts in the appropriate fields. They would, from time to time, hold workshops, seminars, or conferences on particular topics. At least some of those events would be public, live or online, to help build public knowledge of, and attention to, these technologies.

The professionals would hold discussions among themselves to prepare a preliminary list of technologies of interest that they would then discuss, publicly and perhaps privately as well, with one or more advisory groups. Their annual report would examine the chosen technologies or scientific issues based on their plausibility, likely time frame, and potential impact. For each targeted technology, the report would also include fairly specific suggestions for appropriate actions, ranging from further scientific research to public discussion to consideration of regulatory options. This report would be released publicly for discussion and debate.

And then the group would move on to the report for the following year. Reports in the succeeding years would both report on any technologies that newly fell into the areas of interest and update when appropriate the group's assessment of the technologies previously, including its recommendations for possible action. The group would be subject to periodic reviews and decisions whether to extend its mission in, say, 5 year increments, could be required 2 or 3 years before the end of a term. It could probably function with a budget of five to ten million dollars a year. (A larger budget could, of course, support a large institution.)

The Horizon Scanning Group would have to be highly visible, with enough expertise to claim attention, preferably throughout the world. In addition, it would have to make an effort to disseminate its findings widely—through the media (traditional and new); to the relevant policy makers, in and out of governments; to foundations and other non-governmental groups; and to the interested general public.

Perhaps it should release and promote something like a "top ten" or "top five" list at the beginning of each new year, along with substantial background information for those interested in more detail. These discussions in these lists would also make recommendations for which technologies need deeper study, and how soon that study will be needed.

3 Technology auditing groups

Too often, attention to changes—in policy, in law, in society—ends when the changes are adopted. A demand arose that something be done, something is done, and attention shifts elsewhere. Adopting a change says little about what its effects will be. New initiatives often fail; sometimes they are eliminated, sometimes they continue for years or decades continuing to achieve little or nothing beneficial. However, the politicians who pushed for the adoption of a change will take credit for it, but then ignore it. Their incentives to make sure the change actually works, in terms of their chances of re-election, are minimal; by the time the changes results are clear (if ever), the political discussion will have moved on (Mayhew 2004).

This will be particularly true of new technologies. Even the best attempts at foresight often fail. Who would have believed that automobiles would cause lung disease, heat the planet, and change romance—let alone produce the drive-in movies? From the first use in 1945 of the ENIAC, the first electronic, programmable, general purpose computer, to the Internet to the smart phone would have been an unbelievable projection. In addition, from Crick and Watson's double helix paper to forensic databases, direct to consumer genetic tests, and non-browning apple slices would have been a long, long stretch.

One of the things that makes technological advance hard to forecast accurately is that technologies interact. The smart phone required both powerful miniaturized computers and cellular phone service. Transplanting a genetically modified pig heart into a person required genetic technologies but also substantial advances in transplant techniques as well as in assisted reproduction (at least in pigs).

As we cannot predict the future of even launched technologies with confidence, we need to monitor that future, to watch carefully the technologies' effects, existing and nascent. We need the equivalent of a technology auditor.

3.1 Can technology auditing be worthwhile?

Corporations and many government agencies are subject to regular audits. One part of an audit task, performed by an outside auditing firm that has its own reputation—and potential liability—at stake, is to verify that the audited body actually had the costs and revenues it claimed, but audits can also see other effects the firm's actions may have had.

Some governments have bodies that serve similar auditing functions, for financial probity and beyond. In the U.S. the Government Accountability Office (formerly known as the General Accounting Office) will, at the request of Congress, look deeply into particular departments or programs. Many branches of the U.S. government have Offices of the Inspector General. These investigate potential scandals or abuses but can also examine into particular programs. The GAO is independent of the executive branch and the Presidential administration. The Inspectors General are nested inside departments but by both statute and custom, they have substantial independence from the leaders of those departments.

Sometimes when new programs are created in the U.S. federal government, Congress will require a report on the program's progress from the department that implements it. However, these reports typically do not come from an independent entity but from the department that is charged with implementing it. In addition, although these reports often have deadlines attached to them, those deadlines are regularly ignored, sometimes forever.

I know of no "auditing requirement" for a new technology, particularly for the technology itself, independent of a particular statute or governmental program. The closest thing I know to this is the idea of independent "Observatories."

For example, as CRISPR rose to prominence, an international group of academics announced that they were setting up a "Global Observatory for Genome Editing." This idea grew from a conference held at Harvard University in April 2017 called "Editorial Aspirations: Human integrity at the Frontiers of Biology." That led to the formal proposal of the Observatory in an article in Nature in March 2018, 8 months before the CRISPR babies scandal (Jasanoff and Hurlbut 2018). The Observatory's plans were discussed in more detail in a two part, 30 author article in June and July 2018 in the journal Trends in Biotechnology (Hurlburt 2018; Saha 2018).

The Observatory's webpage say it was established just over 2 years later, in September 2020 (Global Observatory). (I do not know but it seems at least plausible that the publicity from the He Jiankui affair helped move it from an idea to a reality.) It is directed by Professor Sheila Jasanoff, one of the world's most noted scholar of science and technology studies, and a professor at Harvard University's Kennedy School. It has two co-directors, Professors Benjamin Hurlbut at Arizona State University and Krishanu Saha at the University of Wisconsin-Madison, and two postdoctoral fellows, as well as a distinguished advisory board. Its webpage details several publications, events, and "interventions" the Observatory has completed. No information seems to be available on the Observatory's budget or on the source of its funds. From its size and activities to date, I suspect its financial resources are not large. It is unclear whether it reports to anyone.

Has the Global Observatory for Genome Editing proven the idea's worth? Not as far as I can tell, at least not yet in its roughly 2 year history. However, it is at least an example that the idea of an observing, if not quite auditing, function for technology has some support.

3.2 Can we create useful ways to carry out such a technology auditing function?

I propose the establishment of specific organizations, similar to this Global Observatory but with clearly spelled out goals, procedures, and reporting responsibilities to an umbrella organization with oversight responsibility for all of the specific groups. Technology Auditing Groups should be similar to, but somewhat different from, Horizon Scanning Groups. Like the Horizon Scanning Groups, they will need some continuity of mission and membership over time and broad credibility. In addition, they need to be small enough to be nimble and affordable.

But the overall structure would need to be different. Tracking and assessing specific technologies, in their manifold uses, across the entire world will require more groups, with more focus—and more funding—than scanning the horizon for plausible new and important technologies. An overall organization would choose technologies to be audited, and then create and oversee several Technology Audit Groups, one for each technology to be followed.

Technology Audit Groups would have fairly short terms, renewable if the oversight thinks continuing them is useful. Each Technology Audit Group would likely need as many professionals as the entire Horizon Scanning Group—as many as ten professionals with the relevant expertise, along with support staff.

Each Technology Audit Group would be charged produce a regular report on its technology. Although the exact contents of the report will, no doubt, vary with technologies and evolve with experience, they might have four parts: (1) the technology's changes and developments, (2) its patterns of use around the world, (3) its benefits and risks as thus far perceptible, and (4) recommendations for further action. Like the Horizon Scanning Group, Technology Audit Groups would base their reports on their own investigations as supplemented by meetings, some public, on particular topics. The Technology Audit Groups would make a draft of their work available for comments—from the industries involved, governments, non-governmental organizations, and individuals—before finalizing the report.

The reports would be issued on a regular and pre-announced basis, though whether that would be annual or not may well depend on the pace of the technology's spread and its consequences. The umbrella entity would determine the timing of the reports and how long the Technology Audit Group should continue.

To give some specific examples, if I were to propose some subjects for audit groups today, I might suggest these five: human genome editing, artificial intelligence, utilization of outer space, new energy technologies, and changes in agriculture. These may not be the right five; another ten or more quickly come to my mind. Which ones would get Technology Audit Groups would be a difficult decision for the umbrella organization, based not just on their importance but on the feasibility of a technology audit and the existence (and quality) of other efforts to assess them. Many projects, large and small, seek to follow these and other new technologies. The Audit Groups, should offer special value through clear expertise, a global reach, and a growing reputation for impartiality.

Of course, the results of such efforts are only "useful" if they are, in fact, used. Many academics, governments, and businesses (including financial organizations from venture capital to insurers to pension funds) would likely welcome the facts laid out in the reports. We cannot know whether the policy recommendations in them would be adopted or not; neither can we know whether they should be adopted. However, a well-resourced, professional, and independent look at the course and consequences of new technologies may well play a helpful role as individuals, businesses, governments, and others try to make sense of the shifting and erratic tides of new technologies that will continue to wash over our world.

4 Organizational complexities

Thus far, I have given these proposed organizations—the Horizon Scanning Group and the Technology Audit Groups—no creator, overseer, or source of funds. These groups need to have some kind of “Oversight Bodies”, institutions that are respected for their integrity and technical understanding. These Oversight Bodies should have long terms—at least 10 years and more likely 20—with the possibility of renewal. However, they pose three difficult organization questions: whether the organizations should be governmental or private, whether the overseeing bodies should be global or national/regional, and how they should be funded. The issues are largely the same for Scanning the Horizon Groups and Technology Audit Groups (indeed, the same organization could create and oversee each) and this section will treat each in turn before examining two possibly suggestive precedents.

If the Oversight Bodies were created, or largely controlled by, governments, they would risk being overly influenced by the political interests—military, economic, or cultural—of their creating government. As government bodies, they might (or might not) have increased credibility within their own country but, for all but the most neutral and inoffensive countries, many would face decreased credibility in some other countries. In addition, governments often come with stultifying bureaucracies and political uncertainties. An Oversight Body created when one party controls a government might well fall victim to abolition or, even if contractually protected from destruction, a “withering away” when control changes.

Private organizations may (or may not) have less bureaucracy and more flexibility. They may both be, and be perceived as, less politically biased and more neutral. For profit firms would, of course, raise huge concerns about favoring their private interests over the common good (or the truth). Associations of scientists or respected foundations might provide a good combination of expertise, impartiality, and prestige, although the neutrality of even such respected groups as the Royal Society, the U.S. National Academies, and the Chinese Academy of Sciences would be suspected by some.

The geographical scope of the organizations is also tricky. The effects of important technologies will not be confined to one country or even one continent. In addition, an organization with a narrow focus—set up by the United States, the European Union, China, or another jurisdiction—is likely to have a limited diversity of views. Some problems will look much different—bigger or smaller—in, say, the United States, India, and Kiribati. Effects of a technology on small farmers, for example, would likely get different attention in different jurisdictions. In addition, in addition, and related, to the losses in diverse perspectives, a narrowly constituted body would have less credibility outside its home jurisdiction.

On the other hand, an international organization may be harder to organize and to run. Dealing with politics in one jurisdiction may well be easier than trying to balance the political demands from many member nations. The United Nations

Organization contains both some encouraging and discouraging examples of dealing with a large array of sovereign countries. The insistence of different member states of focusing solely on the problems that most concern them, could be paralyzing. In addition, international bureaucracies might be even harder to navigate than national ones.

Some of these problems might be overcome using one or more non-governmental (or, in the case of the United Nations and its bodies, non-quasi-governmental) groups as either the Oversight Bodies or their creators. A consortium of national or regional scientific academies, a widely active and respected charity, or, perhaps, a component of the United Nations might be an attractive alternative.

Funding will also be a problem. Depending on the home for the organizations, the amount of money should not be prohibitive: five to ten million dollars a year for the Horizon Scanning Group is a small grant for many governments, foundations, or private charities. The Technology Audit Groups' financial demands, about five times greater, would be harder but should be feasible. One problem might be getting the long term commitments needed; governments may be unwilling or unable to guarantee money over 10 or 20 years and the cumulative amounts might strain some private funders, though, again, this should be surmountable, especially if a consortium of organizations provided the money. Money from for profit sources, though, probably should be avoided for (appropriate) fears of conflicts of interest. Some conflicts may exist with governments, foundations, or charities, but not as blatant as those from commercial firms.

Funding, and governance generally, does raise some special problems for Technology Audit Groups. As the Collingridge Dilemma sets out, a technology that is already being adopted will have already led to more special interests, becoming increasingly vested, in its expansion or restriction. Those interests may well have substantial political power, so the umbrella entity will need a powerful position as well as sources of substantial funding that are not at great risk of being terminated or conditioned by interested parties.

Genome editing has led to two large international organizations (other than the Observatory) that may provide somewhat useful precedents. One was led by national scientific academies and one by the World Health Organization (WHO).

Starting in 2015, and driven by the U.S. National Academies' project mentioned above, a group of scientific academies sponsored the First and Second International Human Genome Editing Summits, in December 2015 and November 2018. The U.S. National Academies of Sciences and Medicine, the U.K. Royal Society, and the Chinese Academy of Sciences sponsored the first meeting; the Hong Kong Academy of Sciences replaced Chinese Academy of Sciences for the second.

After that revelation, on May 22, 2019, the U.S. and U.K. groups created the "International Commission of the Clinical Use of Human Germline Genome Editing" (U.S.National Academies of Science, Engineering, and Medicine, 2019; U.S.National Academies of Science, Engineering, and Medicine, Int'l Comm, 2022). The U.S. Sciences and Medicine Academies and the Royal Society serve as the Commission's secretariat. The 18 member commission includes members, almost all academics, from ten countries from around the world. The Commission is subject to an International Oversight Board with fourteen members, representing not just the Royal Society and the two U.S. National Academies, but a global range of groups. Six listed "sponsors," all from the U.S. or the U.K. (three associated with the N.I.H.) have provided the money. The Commission's main product has been a report released in September 2020, entitled Hereditable Human Genome Editing.

WHO created an "Expert Advisory Committee on Developing Global Standards for Governance and Oversight of Human Genome Editing" in December 2018, immediately after the Second International Human Genome Editing Summit and the He Jiankui revelations (WHO, 2018). The Committee comprised 18 members, the vast majority academics, from 15 different countries. The two co-chairs were from the U.S. and South Africa. The Committee seems to have been funded entirely by the Wellcome Foundation, a British charitable foundation. The Committee's main products were two reports issued in July 2021, a set of recommendations for human genome editing and a proposed governance framework for such work (WHO, 2021a, 2021b).[4]

I am not suggesting these as models for the Oversight Bodies. For one thing, these are set up to consider after-the-fact responses, not to look for new issues. They rely heavily on prominent individuals volunteering some of their time, guided by full time (but relatively unknown) staff. In addition, it is unclear how long, and with how much consistency of membership, they will continue. However, they do provide two examples of largely international bodies sponsored by respected organizations with largely non-controversial funding.

One or more consortia of national science academies or an international organization (if a United Nations group, probably only the United Nations Education, Scientific, and Cultural Organization has a broad enough remit), with funding provided by interested foundations and charities could establish the Horizon Scanning Group and create a body to create and oversee the Technology Audit Groups. This would not be a perfect solution; perfect solutions to creating and governing a universally respected, expert, and impartial body almost certainly do not exist—and if they did exist, they would, no doubt, not be politically feasible. However, these ideas may prove useful.

5 Conclusions

Governing new technologies will never be perfect. Part of that will be because of the inherent uncertainty of the effects of new technologies. Part of it will be the result of the inevitable imperfections in any governance. Even if we had Trafalmadorian insights into the exact details of the future, along with a non-Trafalmadorian belief that we could affect those futures, governance would be controversial, because not everyone would agree on what constitute good or bad results.

[4] The WHO group also issued a report called Current Capabilities for Human Genome Editing in May 2021 and, in February 2021, a report on the possibility of establishing a "whistleblowing" mechanism to report unethical science (Perrin, 2021).

But governance of anything—especially new technologies—cannot just happen at one point in time—it must be a process. It needs to start before a policy is debated, continue through its early testing, peak at the time of its adoption, and extend through its implementation. In technology policy we focus on the second and especially third stages, the most visible and the most dramatic moments. We should be able to do better by attending to the times before and after the two set piece moments in the middle. I have no illusion that what I've suggested will, or even necessarily *should*, in this form, be adopted. I do hope this article provides some useful ideas—and stimulation—toward ways we might improve how we deal with new technologies.

Acknowledgements The author would like to acknowledge, and thank, his excellent research assistant, Cassidy Amber Pomeroy-Carter; two anonymous reviewers; the editors of the journal; and all of the participants in the Workshop on Global Governance of Emerging Technologies, hosted by Fudan University of June 23/24, 2022, especially Professor Li Tang for her role as discussant of this paper.

Declarations

Conflict of interest The author has no competing interests, financial or non-financial, related to this article and received no funding in connection with it.

References

Amanatidou, E., Butter, M., Carabias, V., Könnölä, T., Leis, M., Saritas, O., Schaper-Rinkel, P., & van Rij, V. (2012). On concepts and methods in horizon scanning: Lessons from initiating policy dialogues on emerging issues. *Science and Public Policy, 39*(2), 208–221.

Banta, D. (2009). What is technology assessment? *International Journal of Technology Assessment in Health Care, 25*(Supp. 1), 7–9. https://doi.org/10.1017/S0266462309090333

Collingridge, D. (1980). *The social control of technology*. Martin's Press.

Cuhls, K. E. (2019). Horizon scanning in foresight—Why horizon scanning is only a part of the game. *Futures & Foresight Science, 2*, e23. https://doi.org/10.1002/ffo2.23

Fraunhöfer Institute for Systems and Innovation Research. (2015). Models of horizon scanning: How to integrate horizon scanning into European research and innovation strategies. Available at https://www.isi.fraunhofer.de/content/dam/isi/dokumente/ccv/2015/Models-of-Horizon-Scanning.pdf. Accessed 16 July 2022

Gates, W. H. (1995). Horizon scanning: Opportunities technology will bring by 2005. *Journal of Business Strategy, 16*(1), 19–21. https://doi.org/10.1108/eb039676

Global Observatory for Genome Editing, https://global-observatory.org/#. Accessed 16 July 2022a

Greely, H. T. (2021). *CRISPR people: The science and ethics of editing people*. MIT Press.

Hines, P., Yu, L. H., Guy, R. H., Brand, A., & Papaluca-Amati, M. (2019). Scanning the horizon: A systematic literature review of methodologies. *British Medical Journal Open, 9*, e026764. https://doi.org/10.1136/bmjopen-2018-026764

Hurlbut, J. B., et al. (2018). Building capacity for a global genome editing observatory: Conceptual challenges. *Trends in Biotechnology, 36*(7), 639–641. https://doi.org/10.1016/j.tibtech.2018.04.009

Huxley, A. (1932). *Brave new world, a novel*. Chatto & Windus.

https://www.nationalacademies.org/news/2019/05/new-international-commission-launched-on-clinical-use-of-heritable-human-genome-editing. Accessed 16 July 2022b

Jasanoff, S., & Hurlbut, J. B. (2018). A global observatory for gene editing. *Nature, 555*, 435–437. https://doi.org/10.1038/d41586-018-03270-w

Jinek, M., Chylinski, K., Fonfara, I., Hauer, M., Doudna, J. A., & Charpentier, E. (2012). A programmable dual-RNA-guided DNA endonuclease in adaptive bacterial immunity. *Science, 337*, 816–821. https://doi.org/10.1126/science.1225829

Liang, P., et al. (2015). CRISPR/Cas9-mediated gene editing in human tripronuclear zygotes. *Protein & Cell, 6*, 363–372. https://doi.org/10.1007/s13238-015-0153-5

Mayhew, D. (2004). *Congress: The electoral connection* (2nd ed.). Yale Univ.

Nuffield Council. (2022). What's on the Horizon? Available at https://www.nuffieldbioethics.org/what-we-do/horizon-scanning. Accessed 16 July 2022

Saha, K., et al. (2018). Building capacity for a global genome editing observatory: Institutional design. *Trends in Biotechnology, 36*(8), 741–743. https://doi.org/10.1016/j.tibtech.2018.04.008

US Congress, House of Representatives, Committee on Science and Astronautics. (1967). Technology assessment. Statement of Emilio Daddario, Chairman, Subcommittee on Science, Research and Development, to the 90th Congress, 1st Session

U.S. National Academies of Sciences, Engineering, and Medicine. (2020). *Safeguarding the Bioeconomy*. The National Academies Press.

U.S. National Academies of Sciences, Engineering, and Medicine. International Commission on Clinical Use of Heritable Human Genome Editing. https://www.nationalacademies.org/our-work/international-commission-on-the-clinical-use-of-human-germline-genome-editing#sectionProjectScope. Accessed 16 July 2022

U.S. National Academies of Sciences, Engineering, and Medicine. (2019). New International Commission Launched on the Clinical Use of Human Germline Genome Editing

Vonnegut, K. (1969). *Slaughterhouse-five; or, the children's crusade, a duty-dance with death*. Delacorte Press.

WHO. (2018). Expert Advisory Committee on Developing Global Standards for Governance and Oversight of Human Genome Editing, available athttps://www.who.int/groups/expert-advisory-committee-on-developing-global-standards-for-governance-and-oversight-of-human-genome-editing/about. Accessed 16 July 2022

WHO, Expert Advisory Committee on Developing Global Standards for Governance and Oversight of Human Genome Editing (2021a) Human Genome Editing: A Framework for Governance, available athttps://www.who.int/publications/i/item/9789240030060ut. Accessed 26 Aug 2022.

WHO, Expert Advisory Committee on Developing Global Standards for Governance and Oversight of Human Genome Editing (2021b) Human Genome Editing: Recommendations https://www.who.int/publications/i/item/9789240030381 Accessed 26 Aug 2022.

Wikipedia, Horizon scanning, https://en.wikipedia.org/wiki/Horizon_scanning. Accessed 16 July 2022a

Wikipedia, Office of Technology Assessment, https://en.wikipedia.org/wiki/Office_of_Technology_Assessment. Accessed 16 July 2022b