HOUSE OF LORDS

Science and Technology Committee

Corrected oral evidence: Engineering biology

Tuesday 30 April 2024

11.20 am

Watch the meeting

Members present: Baroness Brown of Cambridge (Chair); Lord Berkeley; Lord Borwick; Lord Drayson; Lord Lucas; Baroness Neuberger; Baroness Neville-Jones; Baroness Northover; Lord Rees of Ludlow; Viscount Stansgate; Lord Strasburger; Lord Wei; Baroness Willis of Summertown; Baroness Young of Old Scone.

Evidence Session No. 5 Heard in Public Questions 41 - 59

Witnesses

I: Piers Millett, Executive Director, International Biosecurity and Biosafety Initiative for Science; Sophie Rose, Senior Biosecurity Policy Adviser, Centre for Long-Term Resilience.

USE OF THE TRANSCRIPT

1. This is a corrected transcript of evidence taken in public and webcast on www.parliamentlive.tv.

Examination of witnesses

Piers Millett and Sophie Rose.

The Chair: I welcome our witnesses to the committee's fifth evidence session in its inquiry into engineering biology. This morning, we will hear from Piers Millett, the executive director of the International Biosecurity and Biosafety Initiative for Science, and Sophie Rose, the senior biosecurity policy adviser at the Centre for Long-Term Resilience. We will kick off with a question from Lord Rees.

Q41 **Lord Rees of Ludlow:** Good morning and welcome. Could you start by briefly introducing yourselves and your work in biosecurity? Can you set out for us—in general, at least—the implications that the developments in engineering biology, which are the topic of our study, will have for biosecurity? In particular, has the picture changed in recent years? Could you introduce us to the topic in that way?

Sophie Rose: Thank you so much for having me here today. As was said, I am the senior biosecurity policy adviser at the Centre for Long-Term Resilience. For your awareness, the Centre for Long-Term Resilience is an independent think tank—it is based here in Whitehall, actually—working to bolster the UK Government's capacity to mitigate extreme risks. We have three policy teams. Mine focuses on biosecurity, but there is another one working on artificial intelligence and another working on risk management.

My work at the CLTR over the last two years has focused largely on how to leverage and govern emerging technologies in order to mitigate and defend against the full spectrum of biological risks. In the last year, it has narrowed in to focus specifically on the threat of deliberate biological misuse in particular. This means that, over the last year, the majority of our work has centred on one of two areas, both of which I expect we will discuss today: synthetic nucleic acid screening, and risks at the intersection of AI and the life sciences. My background is that I hold a graduate degree in infectious disease epidemiology. I used to be a disease modeller. Before that, I worked at the Johns Hopkins Center for Health Security as a biosecurity fellow.

On how advances in engineering biology might have an impact on the biosecurity landscape, this is a double-sided coin. I acknowledge that many of the capabilities in this space will be critical for basic scientific research and will drive amazing advances in health and agriculture. They will be a big part of the burgeoning bioeconomy. I would also say that some of those capabilities have

the capacity to be misused and cause quite a lot of harm, and that some of them might warrant additional oversight. I will leave it at that level, and we can dive into it later.

Piers Millett: I spent a good percentage of my career working for the United Nations on the Biological Weapons Convention. Since leaving there, I have set up a UK-based biosecurity consultancy company. We have worked for Governments, industry groups, international organisations and iGEM, the world's largest synthetic biology competition. Most recently, I have taken over setting up IBBIS, and we will launch the common mechanism, a gene synthesis screening tool, next week. I will talk about that a bit later.

On the importance of engineering biology, we are beginning to see a genuine engineering discipline for biology. We have talked about genetic engineering but there is very little actual engineering in genetic engineering. We are now seeing engineering approaches being applied to biology. As a result, we are able to make things with biology quicker, faster, cheaper and more reliably. A good real-world example of this is the mRNA vaccines during the pandemic.

This is likely to be compounded right now. Recent developments in AI, machine learning and automated science are only making it quicker, faster, cheaper and easier to make things with biology, regardless of whether that is something to do good or something to do harm.

Q42 **Lord Rees of Ludlow:** These are all very exciting developments, obviously, but they are also scary. Could you address the risks? There are two kinds, I suppose. One is unintentional leakage of pathogens from labs, which has happened even in the UK, from Pirbright and all that. The other, of course, is the possibility of bad actors who have access to labs that can generate dangerous viruses using gain of function, et cetera. How can we make ourselves safe against those without some international agreements and without the intrusive surveillance of large numbers of people? It seems really scary to me.

Sophie Rose: One helpful way to consider this is as a broad risk landscape where different threats can emerge depending on the kind of threat actor that we are talking about and the sort of event that we are thinking about. One way of conceptualising how to mitigate that spectrum is to take a step back and inventory the different kinds of threat actors and the different pathways that you are describing, as well as talking about options for mitigations across that spectrum.

Take, for example, the leakage aspect that you mentioned. The Health and Safety Executive—the HSE—has existing responsibilities here. There is immense benefit in that organisation bolstering a bit more around what it does with respect to reporting on this kind of issue. My current understanding is that HSE goes and evaluates labs of a particular standard; the amount of time and frequency is different depending on the level of security that the lab has. However, I think there would probably be immense benefit in facilitating more transparent reporting on accidents, near misses and safety incidents that take place at all of those facilities.

Secondly, on auditing those findings annually, taking an evidence-based approach to the lab area of this could be really beneficial. I think there is an imperfect understanding of what ends up causing accidents like the ones you are referencing. There is a way to harness existing systems and do better in that respect.

When it comes to the deliberate misuse of these technologies, there will need to be a broad swath of approaches across the engineering biology domain. Some elements of that will look similar, I think. Let us consider the purchasing of synthetic nucleic acids. When we think about how best to safeguard these practices, we tend to think about two things: understanding what the person is ordering and whether we think it is appropriate for them to have it; and screening the customer themselves. We try to understand who this person is. We ask, "Do they have a legitimate reason for wanting to purchase whatever this is? Are they going to use it in a safe environment—that is, in a scientific institution with an appropriate level of laboratory safety commensurate with what they are trying to order?"

Taking that sort of similar approach potentially across the engineering biology domain would produce a better understanding of who are the people who are doing this research, where they are doing it, what their intention is. There are things like that that we could make progress on here.

Lord Rees of Ludlow: If you think that is effective, given international dimensions, do they start somewhere else and bring it in? Is it possible to do these dangerous experiments in rather small labs that are not on your register at all?

Sophie Rose: I will take the first question on how to tackle these problems when a lot of them will need international co-operation. I agree. There are a lot of issues—I guess screening is one of them—where you want good coverage across the international landscape and not just something that is focused specifically on targeting the UK. However, a lot of countries look to places like the United States

and the UK and take their lead on matters related to this. We have seen this recently in AI and machine learning domains with the UK's establishment of the AI Safety Institute. There is probably immense benefit in having some countries get their house in order, so to speak, in order to be able to set the standard for what best practice looks like in some of these spaces. We have seen the US and the UK take the lead on this in different areas. Piers might be better positioned to comment on your second question.

Lord Rees of Ludlow: Piers, would you like to comment briefly on the second question?

Piers Millett: Certainly. I think the international aspect here is that there are huge opportunities to embed this in the broader discussions on the bioeconomy, because there are countries in interesting places around the world that are investing heavily in making things with biology. Lots of those capabilities could be misused to cause deliberate harm. Some parts of the world do not have the UK's long history in the oversight and regulation of emerging dual-use technologies. There is definitely a capability out of small labs. You do not need a maximum containment lab to be able to do something really quite nasty.

Anecdotally, when the UK had an offensive biological weapons programme during the 1940s, we did not operate any facility that would come close to a BSL 3 lab, let alone a BSL 4 lab these days. So with good practice it is quite capable of using old technologies to do quite a lot of harm. Using the new capabilities, that is only compounded.

Lord Rees of Ludlow: Any solutions?

***Piers Millett*:** Absolutely. Picking up on what Sophie was talking about, I would double down on the importance of gene synthesis screening. Within the last 10 days, the ISO has just put out a new international standard on nucleic acid synthesis—ISO 20688-2:2024—which outlines good practice in making nucleic acids. It includes a short section on safety and security, including screening. There is a huge opportunity here for a sub-standard that we can work on collaboratively to establish good industry practice. We can build off the work that is already being done, and complement that with additional research to understand where scientists around the world are buying their DNA. At present, there are some important information gaps in understanding who is making it, who is reselling it and who is screening for customers and orders, as Sophie outlined.

The Chair: Perhaps we could leave a deeper discussion to a later

question, if that is okay.

Baroness Willis of Summertown: Thank you. It was slightly terrifying to hear your replies on that previous question. I want to move on to the existing regulatory framework and policy efforts like the biosecurity strategy, published last year. How good do you think the framework and policy are at addressing biosecurity risks from the prospect of engineering biology? You are both, I believe, on the newly established Biosecurity Leadership Council. Is that effective? What have you done so far that you would say that this has made the world or at least the UK a safer place? ***Sophie Rose:*** It is worth saying that we really commend the Cabinet Office and all those who contributed to the 2023 biological security strategy. It was an immense improvement over the previous strategy and laid out some truly vital and important commitments. Three elements in particular are worth highlighting. Updates to the Government's existing biosecurity governance structure include formalising central leadership and providing more opportunities for those people to come together and report on progress, which is valuable. They made some excellent commitments to leading internationally on responsible innovation, which is closely related to what we are discussing today. The strategy made a commitment to investing in the UK's capabilities with respect to bio surveillance. All three of those elements were things that we were really excited to see.

Are they effective? Like many things, the value of that strategy will be in its implementation and which aspects of it are ultimately realised. On that point, although the Government provided a summary of short, medium and longer-term priorities for the implementation, which was great to see, we would be keen to see more specific sets of milestones and targets for each of the priority outcomes in the BSS.

Baroness Willis of Summertown: Could you give an example of a target that you would like to see or a milestone that would be good to have in place?

Sophie Rose: The outcome relating to responsible innovation talks about it in a broad sense. I guess you could see one target milestone as, by mid to late 2024, an articulation of the areas in which the UK is intending to serve as a leader on responsible innovation, and a plan or a road map for what the Government would like to achieve by the end of the strategies period.

The second thing that we are keen to see is sustained resourcing that is commensurate with the risks that are outlined in that strategy. It is crucial for government departments to be able to

deliver on the commitments that are laid out there. Many of us in the biosecurity community see that sort of funding as an insurance policy against needing to fork out in the face of things like Covid-19. There is a great amount of importance in seeing commensurate funding to back up the commitments that are made whereby the departments are resourced to deliver on them.

Piers Millett: I agree with everything that Sophie has just said. As a second example—not to harp on too much about gene synthesis work—but there are the commitments in there to re-examine this issue. Again, it would be useful to see milestones and deadlines specific to that built in supplementary to the strategy.

Baroness Willis of Summertown: Can you give an example of a milestone or outcome that you would like to see?

Piers Millett: Absolutely. The US executive order on AI included 180-day limits on coming up with national frameworks and for the national standards body to start consultations with relevant actors on what these standards internationally, or specifically domestically in the US, could look like. We could see very similar things here. As Sophie said, we could see time limits of six months saying, "We should have started work on this. There should be a process. We should start getting input from relevant actors on how to turn this commitment into action".

If it is useful, I am happy to say a few words about the Biosecurity Leadership Council. One underrated value of the council is providing a link between those doing the work and those in government and in security settings who need it. Operationally, if a scientist or researcher comes across something that makes them uncomfortable, how do they communicate that to the Government? Prior to the creation of the Biosecurity Leadership Council, there was an anonymous tip line that you could contact the Metropolitan Police, but the threshold for using that was quite high. A more realistic scenario would be, "Something's gone on in my lab", or, "I'm aware of this other researcher who's asking awkward questions".

Having a more informal interaction and interface between the two communities where discussions can happen and information can flow so there is no barrier to providing information, especially from researchers into government, is vital. It is part of the Tripwire initiative that the US implemented using the FBI. That does not scale to the UK.

The Biosecurity Leadership Council provides a standing body with trusted people on both sides who can straddle that interface. In

practical terms, we have been providing advice on gene synthesis screening, including explicitly on how to balance some of the security control elements with innovation. It is a really good example of the broader shift to go beyond straightforward regulation to governance. How do you change cultures in communities? How do you engage all the relevant stakeholders and have co-development of ideas and solutions? I see the Biosecurity Leadership Council as a vital step in doing that.

Q43 **Baroness Willis of Summertown:** You mentioned the requirement for long-term resource, but where is the greatest need for the resource?

Piers Millett: It is more in the flexible nature of the resources. A concrete example between the old biosecurity strategy prior and the updated one is that the funds, the budget, that the UK was using for co-operative threat reduction around the world—engaging with scientists and stopping them doing nasty things, or engaging with those who have the capabilities to cause harm and making sure that they are not doing nasty things—basically disappeared between the two biosecurity strategies. My understanding is that there are now funds available and that CTR work is back on in the Ministry of Defence. However, it is open to the vagaries of future budget setting as to whether it will be there next year, the year after, or the year after that. It is a question of sustainability and how to address that, rather than individual additional resources at this point in time.

Baroness Willis of Summertown: Is the Ministry of Defence the right place for that budget to sit for engineering biology?

Piers Millett: For the co-operative threat reduction part? Yes. Sophie mentioned responsible innovation. I understand that you will be speaking to those folks at some point. They have an incredible role to play in the responsible innovation part of engineering biology. That is different to the co-operative threat reduction part.

Q44 **Viscount Stansgate:** How plausible do you think some of the scenarios of concern are to everybody? For example, how plausible would it be for non-state actors to engineer a bioweapon? To what extent would using AI technology make it easier to do so? What critical points of intervention in the development of such would make this more difficult to accomplish? That is just my opening question.

Sophie Rose: I look forward to the rest. It is an important question. One thing I wanted to say in response is that the term "non-state actors" is a very broad group. There is some value in

getting into the nuance. That term, as it stands, refers to groups and individuals very widely in their biological capabilities, their access to resources, their underlying motivations and their objectives. What a well-resourced terrorist group or highly scientifically skilled individual might be capable of will look very different to what an average individual will be able to achieve, whether or not they have access to engineering biological capabilities. If we are talking about engineering a bioweapon at the level of enhancing the function of a known pathogen—let us say, taking something that we have already seen in the world, but making it more likely to evade human immunity or a vaccine—that is a very sophisticated capability. There are advances in technology that make aspects of that capability easier.

You brought up AI, which is a great example. There are AI-enabled tools, biological tools, that are specialised when it comes to different biological capabilities. As they are currently, they would allow you to do things like identify a set of viral mutations that makes that virus more likely to evade human immunity but is just providing one piece of enabling information. As you almost imply in your question, actors need to be able to successfully execute on an entire pathway of things in order to be able to sort of bring about harm in the world.

So achieving that is not impossible, but when we are talking about particular sets of engineering capabilities, it is worth saying up front that that is a sophisticated capability. I would not expect to see an undergraduate trained in biology, for example, capable of doing that at the current stage. I would say, though, that those capabilities are becoming increasingly accessible. Again, if we pivot back to AI and machine learning for a moment, some of you might be aware that Rand and OpenAI have conducted excellent studies looking at whether AI is providing what we refer to as uplift for people who are seeking to cause harm in biology. They ran a set of studies in this space and ultimately found that it did not appear to be particularly helpful in the planning stage of an attack.

That is just one aspect of the pipeline. One way in which AI might enable people in this space is its ability to provide specific live troubleshooting as people are doing scientific experimentation. For those of you who have done wet lab research before, it can be very frustrating when you are tinkering around at the bench, your cells are not growing and you do not totally understand why. Often, in that sort of setting, you go to someone like a postdoc or somebody more senior than you in your lab to try to help you figure out what on earth is going on with that protocol. That can stop quite a lot of people from being able to make progress throughout that pipeline.

You can envision a world in which you can have an ongoing conversation with an AI chatbot to help you to understand and troubleshoot that. That is one way in which that technology is serving as an enabler.

You asked about the clear points at which we can intervene. Two came to mind. The first is access to materials. We think about this broadly in the CBRN space in general. The screening that Piers is referring to is one mitigation that relates to controlling biological material in this way. Another aspect that has opened up as AI has become more of a topic of conversation is trying to consider whether there are certain questions that you might not want these models to answer, or certain domains of information that you might not want those models to provide answers on. That can get technically complicated for lots of reasons, and sometimes it is hard to discern whether that information is dangerous or not.

In terms of points of intervention, thinking about how to stop actors getting access to what will really enable them or to material that they previously did not have before, there are two places to start here.

Viscount Stansgate: What are the surveillance and monitoring systems currently in place to monitor possible misuse, for example, including anything that might identify that someone's wish to access certain materials raises danger signals?

Sophie Rose: With respect to materials, this is where the nucleic acids screening that we have been talking about comes in. I want to be clear that this is not mandatory for companies. There are companies that are signed up to the International Gene Synthesis Consortium, the IGSC, that commit to following a set of principles relating to this screening, but not all companies have to be part of that. For companies that are, the commitment they make is what we started talking about earlier. They are doing some sort of sequence screening, looking at what their customers are ordering, and deciding whether they think that is appropriate. They are screening their customers to understand who this person is who is trying to access this material and, if they are trying to access something dangerous, their specified reason for wanting to do so.

It is important to say that there are no great shared definitions of the things that I am talking about. When it comes to ordering something dangerous, we usually refer to something called a sequence of concern, an SOC. There is no clearly specified definition of what exactly that entails. Companies have different opinions on this, because there is no existing gold standard. When it comes to checking whether your customer is legitimate, there is

no agreed set of standards as to what would make a customer or their use of that dangerous item legitimate. That can make screening really difficult for companies. Those are the existing issues.

Q45 **Viscount Stansgate:** Thank you. That is very helpful. Finally, we have heard from biofoundries that there are procedures in place now—they have in effect been referred to—to screen DNA for malicious sequences, but they are voluntary. Is it your view that it would be a huge help for these to be mandatory in some way? What are the challenges in trying to make these screening procedures mandatory rather than voluntary?

Sophie Rose: I will speak to that briefly and then Piers is well positioned to answer this question. Would it be valuable to make these mandatory? Yes, I think so. When we have spoken to companies that provide this as a service, they have also indicated that it would be helpful for them for this to be mandatory. Two reasons often came up. The first is that this levels the playing field. It costs money to perform this screening. Sometimes resolving queries about whether to give someone something can be expensive, because it requires expertise to understand bioinformatically what is being ordered. So it levels the playing field, by making sure it is not a competitive advantage for a company not to screen.

Secondly, it helps to justify the process to their customers. For people who are being asked for information about who they are, where they work and why they want something, it can help for companies to be able to say, "I have to ask you for this information". I have outlined over the course of this conversation some of the reasons why not having great shared definitions for some of these things is challenging.

The Chair: Is there a definition of what a malicious sequence is?

Sophie Rose: Actually, yesterday, the White House delivered on the commitment that Piers referenced earlier and released a framework for its expectation on what providers in the US are expected to do. This included a definition of what they see a sequence of concern being. For them, it is a sequence that matches anything on their federal select agent, or FSAP, list. If the order is being provided internationally, the definition is that a sequence of concern is something that matches what is on the commerce control list or export controls.

They have a vision that over the next two-and-a-bit years, so by October 2026—hopefully my maths is okay—that definition will be

broadened to include more functional definitions, so instead of just having to match something on a list, the definition will shift to knowing that a particular sequence can confer toxicity, pathogenicity or increased transmissibility. It is technically difficult to screen for that, currently. My understanding is that that is why they are currently leading up to that as opposed to requiring it.

Viscount Stansgate: I am interested in what Piers has to add in this area.

Piers Millett: Just to double down, yes, screening requirement is important. To unpick that a little bit in the UK context, there is the potential to require companies that make genetic material, or machines that make genetic material, to screen. My understanding is that the number of companies involved is quite limited in the UK. Perhaps a more relevant part is building into our regulatory systems how we place a requirement for scientists and researchers to use companies that screen so that they are part of doubling down on good practice.

The US has chosen to link that to funding. Given the scale of UK science funding, that might be considered here. To dip into this issue of the databases that sit behind the screening systems, now is a uniquely important time to have a broader international discussion on this. The US is setting its own rules domestically. If we fail to have some discussions, the rest of the world will effectively be left with adopting the US standards or developing something that is potentially not synergistic with their approach.

There is a huge opportunity here. There have been some well documented challenges in getting companies to screen DNA. Some of those is the cost, as Sophie mentioned. Some of them are sensitivities over the orders. Do companies want a third party with access to their order information? That is why we have built a free, publicly available tool that allows anyone, anywhere to screen specifically to remove some of the barriers that have been flagged previously. I think I will leave it there.

Q46 **Baroness Young of Old Scone:** I want to ask about gene drive processes specifically. There has been quite a lot of international coming and going on that aspect of agriculture and whether it is safe. There have been various moratoriums. Is that area sufficiently well monitored where there are international standards? I know that the Royal Society has published a guide to gene drive research, but we do not have anything out there yet in real use. What are your comments on security and safety in that area?

Piers Millett: I was at the big international research meeting for gene drives in the west coast of the US about three weeks ago. What is fascinating is the amount of effort that has gone into building regulatory capacity in places where gene drives may be used—for example, in Africa—being supported by the Wellcome Trust, among others. There is a drive to make sure that the procedures and practices are in place for the safe, effective and responsible use of these technologies, and with the full recognition that they have the potential to move across borders.

I also know that the World Health Organization is in the process of exploring governance frameworks for gene drives, in addition to the excellent work done by the Royal Society in the US national academies. That is with potential partnerships with UK and US-based funders but with a global remit to see how rules and good practices need to be adapted into different settings to work operationally—for example, in south-east Asia as opposed to Africa.

Baroness Young of Old Scone: How big a risk does it represent at the moment?

Piers Millett: From the deliberate misuse side, it is a future risk as opposed to a short-term real risk. As far as the accidental impact of the drives is concerned, that is a probably a bigger issue to consider. I am also aware of the safe genes programme that was run by DARPA in the US, which looked specifically at how, if something went wrong, we could roll back a gene drive. How could we stop its transmission? How could we interfere with a gene drive that had been let out? How could we overwrite something to bring it back to a natural state?

There has been some excellent work done to develop safety technologies associated with gene drives. It is comparable to imagining how the world would have been different if we had had seatbelts from the first day we had automobiles, rather than waiting 50 years to put them in cars.

Q47 **Lord Drayson:** You have already explained how AI machine learning is really an enabling technology that is having a massive effect. I should like you to expand on the implications you see of the adoption of this technology. In doing so, can you explore how it is affecting our ability to monitor and respond to biosecurity threats? Sophie, you already mentioned the potential proliferation of lab expertise via a chatbot. What research is being done to provide guardrails to large language models? What sort of research should government be ensuring is done in this area?

Sophie Rose: It is a great question. This is an area in which I commend both the UK and the US. They have taken a bunch of steps over the past 12 months to start putting themselves in an excellent position to tackle this problem. The UK established the Foundation Model Taskforce , which is now the UK AI Safety Institute. The US has also set up an analogous institute domestically. Those safety institutes have signed a memorandum of understanding, if I understand correctly, such that they can work together on some of these topics, which is an excellent idea.

On how this is impacting our ability to monitor and respond, as well as what kind of research government should be doing, there are a few things here. This is a space that is moving very quickly and is unfamiliar to a lot of us. As a result, there are probably many risks that we do not understand currently. It is not super clear where and how those risks will manifest.

So I see there being a lot of value in building government's muscle in acquiring the appropriate sort of expertise and funding places like AISI (the UK AI Safety Institute) to do some of this work so that they can step back and say, "Okay, what are these models capable of doing? How dangerous do we think that capability is in the hands of different threat actors?" It is about building the muscle to be able to do that and to do it periodically, so that it is not just a single snapshot or point in time; it is about being able to say with confidence, over time, "We're seeing an improvement in this capability. We think that capability has implications for misuse", using that as a temperature check for where the risk is coming from—and, therefore, where mitigations need to be applied.

Taking that sort of approach at this stage is important. I do think that our understanding of these risks is too premature to do anything that looks like regulation currently, but I see this as a space where the risks could ultimately be sufficient to warrant that and to get a bit more specific about the sorts of capabilities that we might want to assess.

It is important here to be able to say how somebody having access to this model does anything different than them having access to the internet, and how in particular it is enabling them. Let me go back to the troubleshooting experimental example. The internet is okay at this. There are forums where you can post and explain, "I'm doing this experiment. This is what's happening. Can somebody please help me?" For obvious reasons, being able to do that live has its advantages as we continue to see things like multimodality be incorporated into AI models. You can also envision being able to do things like giving the models pictures or

readouts of your results in addition to asking them questions. You can start to see where, in terms of how much help they offer, the internet and AI start to diverge.

Having a team of researchers who can ask those questions is really important. However, that requires buy-in from companies, as well as access to be able to ask those questions.

Q48 **Lord Drayson:** To your example, the security services can monitor an internet chat forum to see whether people should be talking about certain things or not, but when it comes to the Government wanting to regulate this space and the security services wanting to contain expertise limited to the use of large language models, how can the Government effectively constrain the companies that have the foundational models when they are not UK companies?

Sophie Rose: That is a good question. First, we want to be really careful about restraining information. A lot of thought needs to be put into what falls into the quadrant of whether this is sufficiently dangerous such that you want to keep that information away from people.

It can be difficult to do this in practice, even before we consider how you get company buy-in, because it is sometimes hard to tell whether the questions that somebody is asking are dangerous. We refer to this as the dual-use research space, but, as you can imagine, up to a point those questions might sound like those of a PhD researcher. If we are talking about misuse, there comes a point in the pipeline where those questions start to diverge from someone doing peaceful work, but it takes a long time and is difficult technically to ask a model to recognise that. We have safeguards for things like bias, saying inappropriate things or generating inappropriate content, but that technique does not work as well for some of these things.

Lord Drayson: Piers, could you expand on that? I am particularly interested in the challenges that Sophie has articulated about the regulation of this technology, given that it is moving so fast and it is hard to determine applications that are good and applications that are evil.

Piers Millett: Absolutely. This is the space with emerging technologies that sits ahead of regulation. You are trying to encourage actors to be good and set up the guard-rails that you were talking about in order to apply some degree of transparency. All that will fall short of a regulatory approach, but it is vital for learning what future regulations might need to be.

Also vital here is that dealing with some of these intangible technologies will continue to be an incredible challenge. As Sophie mentioned, this is only one step on the pathway that a malicious actor would have to go down. We have a much better ability to interdict or interfere at the point where that digital information is translated back to biology. That is why we are seeing this doubling down on nucleic acid synthesis screening and gene synthesis screening, I think. If you have designed something really nasty online, ultimately to cause real-world harm, you need to turn that back into biology. Knowing who is ordering what, using the tools and technologies that we have, and even more importantly, strengthening those screening tools so that they are increasingly resilient to these AI risks, will be important. There is a research agenda there.

Something that has not come up yet but is also vital is thinking about vulnerability reporting. What happens when we know there is a gap in the system? What happens if one of these companies identifies a way to misuse their tool that could be really quite concerning from a biosecurity point of view? I believe that having good rules of the road—about what happens, who talks when and what their thresholds for making information public are—is quite different in the biological space compared with cyber space. It is not necessarily going to be easy just to move the rules and common practices of cyber vulnerability reporting into biotechnology. We will need a process that involves regulators and companies developing what the rules of the road for vulnerability reporting need to be. I do not see that happening so much at the moment.

Q49 **Lord Drayson:** Can I probe you on the point you made about the sense in regulating at the point of the transition from the computer science to the biological science? Are we already at a point where there are certain types of research facility that the Government should be controlling more tightly? Given what you said earlier about the facilities in the 1940s, is this really practical, and should we be doing something about it now?

Piers Millett: This comes back to who is making this synthetic DNA. There are a limited number of companies in the UK. My understanding, or my awareness, is that all the companies in the UK that are making DNA, biofoundries or machines that are able to print DNA are engaged with the discussion about how we can realistically screen customers and orders. So it does not seem to be a huge imposition for the UK.

The real challenge is this: how do you spread this practice all over the world so that a UK researcher cannot just order their synthetic

DNA from somewhere else—or, indeed, so that a malicious actor somewhere else in the world cannot simply circumvent controls in Europe, North America, the UK and a couple of other places and order things that could be used to cause harm comparatively easily? That is why building an international framework around this using the standards that I mentioned—the ISO standards backed up by an expectation in some of the security forums, like the Biological Weapons Convention or UN Security Council Resolution 1540—will provide a stronger framework and allow us to shine the spotlight on those who are not screening. If there are companies out there that are happy to supply anyone with dangerous DNA, we need to know who they are and to remove the excuses for not screening.

Q50 **Lord Drayson:** Can I probe you a little further on the balance that needs to be struck between ensuring responsible innovation in this space and constraining the wrongful use of the technology? In the international context, it is quite striking that the US President's executive order on AI said explicitly—in the first paragraph, I think—that it was about making sure that the US remains the world leader in AI. That is in some contrast to the UK's policy, which did not really make any clear pronouncements about the importance of the technology delivering economic growth.

How do you see the international development of the balance between those two conflicting elements being achieved, and what should the UK's sovereign capability be to ensure that the UK is a leader in responsible innovation but, at the same time, is generating economic growth from the application of engineering biology in the UK?

Piers Millett: We are already heading in the right direction on that one. I believe that the engagement of the engineering biology community, including through the Engineering Biology Leadership Council, has had a long-standing interest and willingness to engage with exactly this question. Where we struggle in some of the other technology areas is getting buy-in from key actors. In the UK, we have the BIA, the big industry body that that has already engaged on this issue. We have been holding academic conferences on that have specifically addressed this issue in the UK for a long time. I think we are doing exactly the right thing.

If it is a question of what more we could do, one area that we have not fully explored is how we use our markets to create a greater incentive for other people, in other parts of the world, to do this. If, for example, you want to access the UK market to sell the product of engineering biology, demonstrating that your research and development, included screening of synthetic DNA, could be a very

minimal requirement. It is part of the paperwork that should already exist and could easily be part of an approvals process that would mean developers around the world will want to do this, because ultimately they want to sell their products in North America, Europe and the UK. The way we use the market to incentivise some of these security issues has the potential to be explored more fully.

Sophie Rose: I would add that although the framing of the US AI Executive Order (EO) is focused on balancing these things in order to have growth, it contains a lot of firm commitments to making sure that the safety is there to allow for that growth. As regards some of the commitments that Piers and I have referenced, them taking on synthetic nucleic acid screening today comes out of that EO, for example. I also commend the way in which the UK has chosen to approach this so far.

Fundamentally, at the end of the day, if you are going to rely on a bioeconomy strategy for economic growth, you need to have public trust and buy-in to what those things are and what they are doing for the country. Even a very minor incident in the space could undermine our ability to do that, both in the engineering biology sector and in the AI sector. Those things go hand in hand. I acknowledge that the framing in the US piece was a bit more tilted towards that.

Q51 **Baroness Neville-Jones:** Perhaps I may turn to resilience for a moment. In a lot of threat landscapes these days, resilience is part of the game. In defence, it is essential but not sufficient. When something happens, you need to be able to respond rapidly. What are the key elements in resilience in this area? In particular, what about the relationship between bioengineering and the public health system generally? That seems to be a component part of resilience, but you may have other thoughts on where we need to build our capability to recover. I am interested to know what you both think about how important and salient in government policy this aspect is.

Sophie Rose: Just to make sure that I understand the question correctly, it is about the other elements in a situation where an event like this occurs and whether they are important to the UK being resilient against those sorts of threats.

Baroness Neville-Jones: Yes

Sophie Rose: I agree with the framing of your question that gestures at how important things like the public health care system are to being be able to sustain this.

Baroness Neville-Jones: The pandemic showed that there is an issue there. Have we learned lessons?

Sophie Rose: Yes and no. There were ways in which the pandemic revealed that planning that had previously been put into place had not considered a broad enough range of possibilities and therefore what actions might need to be taken. It also emphasised that cuts in the public health space can be detrimental to our ability to respond effectively to these kinds of events. It is also worth pointing out at this stage that no matter whether we are talking about a natural pandemic, the result of an accident or the deliberate misuse of biology, once we get to that stage we are playing a very similar game where we are relying on our ability to keep people healthy and safe, irrespective of the origin.

On the contributions to a response, I think that bolstering those elements covers all your possible pathways to harm, as opposed to when you are thinking about like how to prevent these things, where the different mitigations look quite different, whether we are talking about preventing deliberate misuse or accidental things and so on.

Baroness Neville-Jones: Could you argue that the threats in this area are in that difficult quadrant of not very likely but potentially quite serious if they occur? What kind of measures should government take to prepare for that kind of situation?

Sophie Rose: This takes us back to the conversation we had earlier about the biological security strategy. It is an area where it is worth putting effort into prevention. How do you make it harder for people to execute on it? I think about it like the Swiss cheese model, which you may remember with respect to Covid: that the more barriers that are in the way of somebody trying to cause harm in this space, the less likely they are, ultimately, to be able to execute on it. Covid showed us that waiting until we get to the point where we need to respond and then having to do so on the spot is an extremely expensive exercise. In relation to some of the commitments made in that strategy, to me that means that spending at that point feels like an insurance policy for preventing the later on-the-spot spending.

Q52 **Baroness Neville-Jones:** Is the health department involved in strategic planning?

Sophie Rose: Do you mean: is the department is involved in strategic planning and doing preparedness exercises and so on? It depends on the exercise. That is not one of my areas of expertise. Piers, if you are aware of any specifics, please feel free to jump in.

Piers Millett: I am not particularly aware of any specifics, but I would just highlight the well-trodden path of needing to broaden out the set of high-consequence biological risks that we consider. This was a lesson from the pandemic. When we went into it, we focused on influenza, but there is a variety of other nasty pathogens and, increasingly with engineered biology, alternative approaches could have a similar impact.

Lastly, I would highlight some work that Sophie and her centre did on metagenomic sequence screening. One challenge is that we do not know what normal is. We do not understand what pathogens circulate on a day-to-day basis. If we do not know what normal is, we do not know what unusual is, and if we do not know what unusual is, we will find it very difficult to find what unnatural or deliberate is. So in order to be able to see what nasty bugs are circulating at any given time, these screening systems are incredibly important. I am disappointed that we have not continued to invest in those sorts of surveillance capabilities, because they provide many day-to-day benefits on public health but also have critical capability if we move into accidents and, preferably, if we are to avoid deliberate biological attacks. That is also critical there.

Q53 **Baroness Neville-Jones:** What you are saying seems to amount to the need to increase our information base about what constitutes normal in order to understand the abnormal. I come back to a question that was raised earlier. The Ministry of Defence is very much in the centre of things. Is that really the only department that has an interest? It gives a strong national security swing. That seems to be inadequate when it comes to the bio side and the effect on civil society. It seems that there needs to be more players in this game than there are at the moment.

Piers Millett: Yes. This is a fantastic example of how the revised security strategy is so much better than the original version. The drive from the Cabinet Office provides a whole-government approach and allows all the different line agencies to be involved. The shift from the first version to the one that came out last year is incredibly important, as is maintaining the Cabinet Office's role. There are some fantastic diagrams in the biosecurity strategy that highlight how many different agencies are involved and who has the lead in different parts of the national strategy. It covers human health, animal health, the HSE, Foreign Office, MoD and DSIT and, in my opinion, it is an impressive example of an attempt at joined-up thinking in this space.

Baroness Neville-Jones: Is it now a case of implementation? You have drawn up a picture of all the people who need to be involved, but, to your knowledge, is it actually happening?

Piers Millett: To my knowledge, yes, individual agencies and departments that have leads in individual areas are undertaking work. What is less clear to me is the process whereby we sit back and reflect on progress since the strategy came out, what that would look like and how it all comes back together. Having the Cabinet Office in the middle of this makes that so much more practical. The first biosecurity strategy took seven years to develop, I believe. Hopefully, we will not have to wait seven years to get a review and an assessment of what the next steps need to be.

Sophie Rose: The one-year anniversary of the strategy's release is in June this year, and my team is currently working on trying to assess each of the 15 priority outcomes and the three strategic enablers that accompanied those. We are trying to assess the progress on each of those, and I am happy to provide that information when we are finished.

Baroness Neville-Jones: This is all, in a sense, on the threat side. Does opportunity feature adequately?

Sophie Rose: It is not the aim or objective of the biosecurity strategy to maximally consider the benefit side. That is well considered by other government strategies.

Q54 **Lord Berkeley:** How can the situation with the Cabinet Office and the MoD be improved? We have heard a lot of evidence from different people that co-ordination is often lacking. The MoD likes to think that it runs everything, which may be right or wrong—we do not know. However, there needs to be proper co-ordination before you two, who are both experts on many international committees, can use the strategy internationally. Do you have any ideas?

Sophie Rose: This is an excellent question. In any Government, there will always be inherent tensions around these kinds of risks, and many different interests to consider when we are trying to prioritise bioeconomy growth, as we talked about, and how you do responsible innovation well. Of course, that becomes ever more complicated when the risks have implications for national security, as they do here. I want to be cognisant of that tension, and I understand where it comes from, but I do not think we can afford to govern a sector or any kind of technology in silos. This is a space that is moving at the speed of a train, so we cannot afford to be moving at the speed of a horse-drawn carriage when it comes to working together. All these risks are quite cross-cutting in nature, and they will require cross-government ongoing collaboration to

fully tackle them, including co-operation from the national security sector.

Piers Millett: I point to page 59 of the biosecurity strategy, because it nicely lays out how a phenomenal number of different bits of government feed into this.

Lord Berkeley: Thank you.

Q55 **The Chair:** Can you tell us a little about the state of play on international agreements on engineering biology and their biosecurity implications? What could the UK be doing to enhance international co-operation?

Sophie Rose: We have mentioned a few today and have acknowledged that domestic efforts can be very valuable. It is helpful to get your own house in order and understand it before turning to the broader community. Of course, as we have all said, this will require global co-operation.

A few things have been really encouraging in this space. As I mentioned earlier, the joint US-UK partnership on the science of AI safety that was announced in early April this year is a great example of bilateral work. Similarly, the US-UK strategic dialogue on biological security was announced in January. It includes commitments on joint investment in R&D and working together on some of the issues of responsible innovation that we have talked about today, and other areas of investment relating to biosecurity, such as forensic attribution, vaccine development, et cetera. All those things are great.

I would encourage the UK to complement that with additional engagement at the G7 and G20 level. For example, when it comes to things like nucleic acid screening, there is a G7 commitment to an AI code of conduct that includes a bunch of international guiding principles. This opens up the space in a similar way to how the US built that into their AI executive order. You can envision similar conversations taking place at that level. There is good work being done at the bilateral level, although there is room for improvement as well.

Q56 **The Chair:** Are we doing enough to bring countries like China into this dialogue?

Sophie Rose: Piers is well positioned to answer that question.

Piers Millett: With regards to China, it is difficult to single them out. They obviously have huge amounts of relevant technology and expertise. We have been able to work with them in the past, such

as on the study on synthetic biology, as it was then, between the Chinese academy of sciences and engineering and the US and UK versions. It included security components but looked more broadly at how to develop these capabilities.

Certainly at the technical level there are huge opportunities. When we move into multilateral global forums, the broader geopolitics are making progress very difficult. To get substantive work on this into the Biological Weapons Convention will be very challenging in the short to medium term, simply because there are states in the world that do not want to see outcomes from multilateral processes, regardless of the subject matter. We are tied up in that, which is perhaps why we can look at other forums that bring people together.

Sophie mentioned a couple of others; I reiterate the opportunities to do international standards work through the ISO that will help build capabilities in this space. In my opinion, there is quite a lot of work to be done on updating export controls to take into account the realities of engineering biology and synthetic nucleic acids. That is partly domestic, but the UK is also phenomenally well placed to take that to the Australia group, which has some 50 states that harmonise their export control regimes on chemical and biological weapons issues.

Q57 **Lord Berkeley:** Piers, you mentioned China, and earlier I mentioned the MoD. How do you manage the academic links that you need, with all these people worldwide, considering the MoD's sometimes attitude to China?

Piers Millett: I think there is an opportunity for a cultural shift within academia. This is not a new challenge; it is a newer challenge for biology and engineering biology. This is an issue that we have confronted for a very long time in, for example, nuclear physics.

Part of the problem is obviously rules and regulations for locking down information, but perhaps even more important is having the reflection point among the academics—an awareness that their work has broader implications and could be misused and desired by foreign actors—and building that into their day-to-day practices. They should not stop innovation but should reflect on who they speak to and how and where they speak to people about the work they are doing. It is alien to engineering biology and to much of academia, but there are examples.

Q58 **Lord Lucas:** At the end of the day, our report will make recommendations to the next Government about what they should

be doing over their five years to make sure the UK is prepared for eventualities in this area. What are your top three recommendations?

Sophie Rose: The first would be an approach outlined earlier on building the muscle of government to make sure that it is in a position to draw on relevant expertise and have risk assessment processes in place so that it can understand and continue to evaluate risks that are part of the engineering biology landscape. This is good for two reasons. First, where regulation ultimately becomes necessary for mitigating risk, that Government are then well-positioned to move when risk passes the threshold they are comfortable with. Secondly, it allows us to facilitate an environment that is permissive to innovation, because it is trying to ensure that any mitigations in place are commensurate to the risks.

The second recommendation is that the UK Government should take a stepwise approach to building a more robust nucleic acid synthesis landscape. That could start with things like developing public-facing guidance on what best practice in the sequence screening and customer screening elements of those things would look like. What does it look like to do those things well?

With respect to customer screening, we would really encourage the Government to consider where and how those kinds of customer screening practices could extend to other elements of the engineering biology landscape. Where could you use those to help make other elements of this domain safer as well?

Lastly, and we have not talked about this as much today, I think the UK Government should establish a process for better overseeing research that has dual-use implications. We should keep allowing all the benefits that come from that research but, like Piers referenced at the end of his last statement, make people more aware that some of their work might have these implications and ensure that they handle it appropriately to make progress on the problems and the topics they are working on without opening us up to undue harm.

Piers Millett: I would like to see continued progress on requirements to screen, and the requirements to use companies that screen, synthetic nucleic acids. I think there is progress and it would be awful to see the next Government move away from implementing those sorts of measures. They will become increasingly important as the impacts of machine learning and AI continue to play out.

Secondly, picking up Sophie's point on good practice around customer and order screening, I reiterate that there is an opportunity to develop a supplementary ISO standard that would go into more detail on how to do that. It would then become a de facto industry standard in many cases. With larger companies, once there is an ISO standard it is very difficult for compliance officers not to implement it, so I see particular value in that.

Finally, we have the issue of export controls—how we move beyond controlling whole organisms to biological parts and, in particular, databases of sequences of concern, as Sophie said.

Q59 **Lord Lucas:** Thank you. What is necessary to keep your own committee healthy? I have been observing the interactions that you are part of. Is this some organism that will flourish over the next five years, or does it need some help?

Sophie Rose: Do you mean the Biosecurity Leadership Council specifically? I would like to see the current practices continue with a new Government. It is an opportunity to bring forward policy considerations, challenges and proposals and have a discussion among relevant experts. What would be the challenge of implementing this? Will implementing this as it is described capture the risk that we are worried about, or will it unnecessarily capture other things? That level of open back and forth dialogue has been invaluable, and I would urge that it continues.

Piers Millett: As you correctly identified, it is about maintaining what now exists. We have been very lucky that the Biosecurity Leadership Council has been supported by some fantastic folks inside DSIT, and my understanding is that those positions are likely to continue unless there is a major restructuring. That is very good news.

Keeping this issue on the agenda will help keep it as a priority for DSIT. That recognition of the importance of biosecurity writ large, biosecurity as a component of engineering biology, and the drive towards the global bioeconomy will all help to ensure that this continues use into the future.

The Chair: Thank you very much to both our witnesses in this session. You started by perhaps alarming us, but you have talked about some practical and pragmatic ways forward that we should be implementing as fast as possible to try to address this rapidly moving agenda. I think, Sophie, that you said that there was some additional material you could send us. We would be very pleased to receive that as evidence.