



HOUSE OF LORDS

Science and Technology Committee

Corrected oral evidence: Advanced Research and Invention Agency (ARIA) (non-inquiry session)

Tuesday 12 November 2024

10.15 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

Heard in Public

Questions 1 – 26

Witnesses

I: Matt Clifford, Chair, Advanced Research and Invention Agency (ARIA); Ilan Gur, Chief Executive Officer, Advanced Research and Invention Agency (ARIA).

USE OF THE TRANSCRIPT

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



HOUSE OF LORDS

Examination of witnesses

Matt Clifford and Ilan Gur.

Q1 **The Chair:** I welcome Matt Clifford, the chair of ARIA, and Ilan Gur, its CEO, to this evidence session of the committee on the Advanced Research and Invention Agency—ARIA.

We are delighted to have you here. We have been looking forward to hearing about progress with ARIA since before the election was called, so it has been a long wait. The committee is very interested in hearing what you have to say.

The session is being broadcast on parliamentlive.tv. You will receive a full transcript of the session shortly afterwards. Please make any small corrections that you need to. If you think of any information that it would be useful for us to see, please send it to us and we will receive that as formal evidence. If you are settled, we will kick off with the questions.

ARIA was formally established almost two years ago and, over the past year, has announced seven¹ programme directors with major areas of focus. It would be interesting to hear from you an outline of some of the work you have done, what you have achieved so far, and perhaps what some of the successes have been, recognising that what you are doing is a long-term endeavour and this is early days.

Matt Clifford: Thank you, Chair, and noble Lords, for having us today. I am the chair of ARIA. Before we get to where ARIA is now, we both wanted to say thank you all for providing us with this opportunity to speak. We have been looking forward to it for a long time too and there is a lot to update you on. We thank the whole House of Lords for the role it played in establishing ARIA. It was a big and important bet that we are grateful you made.

ARIA's existence is due to a collective effort that began long before Ilan and I were appointed. The whole team is grateful to the parliamentarians who made it possible, and to the Ministers and civil servants who embraced this unusual mission and gave us the mandate, freedoms and independence we need to achieve what we want to achieve over the next decades.

Since we were established in January 2023, we focused on trying to build a strong foundation for the work ahead—a dedicated team, a strong culture, robust operations and a new approach to catalysing

¹Note by the witness: ARIA has announced eight programme directors, leading seven areas of focus (one of ARIA's programmes is co-directed by two programme directors)



HOUSE OF LORDS

scientific breakthroughs. I will let Ilan talk you through the details in a moment, including on how our inaugural cohort of programme directors is progressing. You will see a set of highly ambitious programmes and a set of funding grants that we hope will unlock breakthroughs for the UK and the world.

In my capacity as chair, I see my core responsibilities as maintaining ARIA's level of ambition and risk appetite. That is one of the most important responsibilities of the board in this unusual set-up. In an organisation like ARIA, there will always be a temptation to look for quick wins, to find something to report quickly that looks good and maintain support for it. But that is one of the biggest risks ARIA faces, because if it is to stay true to its mission, the reality is that only a fraction of our programmes will succeed. That is by design, and they will take a long time to do so.

Fortunately, the way in which we have chosen our programmes—you will hear more of this from Ilan—is that even if one of them succeeds, we hope it will change the world. So a big part of our role as the board has been trying to create the right balance in our approach to governance: putting transformative technologies for the benefit of the UK, despite a high risk of failure, right at the heart of what we do, while seeking to uphold our role as responsible stewards of taxpayers' money.

In our approach today, I am happy to field all your questions on ARIA's governance, the shape of our organisation, our responsibilities as an agency and our relationship with government. I will largely defer to Ilan on the details of the programmes, the research approach and how we are choosing to focus our time.

Q2 The Chair: Clearly, we are interested in the fact that you will be taking some big risks, and I am sure that is important. How are you going to decide to stop programmes?

Ilan Gur: That is a wonderful question. The thing that worries us when shaping a programme is not that the programme might fail to meet its objectives, but that it was not designed to have a level of ambition and differentiation where it would not happen without ARIA's intervention. Also, if it were to work, would it really have a chance to be transformative and consequential?

So I think about failure in this way; we use the word "failure" quite a bit in talking about this. As long as we find an appropriate, well-suited, driven, passionate and exceptional group of individuals to fund in pursuit of a target, failing to meet that target will not be a failure of progress. It will be a failure to meet the objectives, but it will certainly drive progress in terms of our learning and advancing



HOUSE OF LORDS

towards the types of inventions and capabilities that ARIA was designed to do.

The Chair: So you are going to go on funding things even when it becomes clear that they cannot meet their targets because they are still doing useful stuff.

Ilan Gur: No. One of the things that Parliament has afforded to ARIA in the way it was set up is to be able to be adaptive in our funding approach. In the programmes that we launched, and hopefully we can chat about this more, we have set these ambitious targets, but our programme directors will be working closely with those teams. Those initial targets in the first steps of that research are very speculative, but not only our hope but our intention is that we will stretch. We may fail along the way, but then we will adapt. Some of those projects certainly will be discontinued if we realise that there is no way to get to the target—the things we are doing are at the edge of the possible—or they may pivot and change directions. That is both the opportunity and the great responsibility that we have.

Matt Clifford: One of the key achievements of the ARIA Act that many of you were involved in was mandating a level of risk appetite that is unusual. There is a potential failure mode for any endeavour like this, which is to disguise failure in order to show a higher rate of success than might be there. ARIA has the statutory mandate to take this very high level of risk in pursuit of extraordinary outcomes, so, from a governance and a culture perspective, we have baked that into everything that we do and not to be afraid of failure. So we will be comfortable if all but one of these early programmes fail to achieve their stated objective, and we take comfort from the fact that that is what Parliament has mandated us to do.

Q3 **Baroness Willis of Summertown:** It all sounds very exciting. Could you discuss the areas of focus that ARIA has chosen, and what was the process that you went through in order to choose those areas? I cannot quite understand from the literature we have been reading who said that these are really good areas, and who then confirmed that those are areas where there is a knowledge gap and a huge potential?

Ilan Gur: Maybe I will start by putting that question into the context of what has happened so far with ARIA. I appreciated Matt's comments, and I echo the thanks to you all for the opportunity to speak here and for the role Parliament played in laying the foundations for ARIA. It is an honour to be here.

I jumped on the opportunity to serve as ARIA's inaugural CEO for a few



HOUSE OF LORDS

reasons. First, I believe in the power of science and technology to drive progress and serve society.

Importantly, the UK is one of the few places in the world from which world-changing breakthroughs can emerge. That is in part because of the powerful base of scientific funding, ideas and institutions that exist here, but above all it is because of the incredible concentration and depth of scientific talent. Funding, ideas and institutions are crucial but, ultimately, it is people who drive innovation. ARIA is a bet that, right now, by activating the incredible base of research talent we have in the UK in new ways, we can transform not just the future of the UK but the world.

Talent sits at the heart of the agency. When I think about our progress to date, one of the things I am most proud of is the team that we have assembled to deliver on the mandate that you have all set forth. If you go to our website, you will see that it says at the top of the page, “empowering scientists and engineers to reach for the edge of the possible”. The way we think about this is that we back people and then projects. I was reflecting on this, and I think this approach started even before I was appointed. Parliament, after a number of debates, decided not to prescribe a narrow focus for ARIA but to entrust the team with this broad mandate to unlock breakthroughs that can be transformative across disciplines and sectors.

We have passed on that idea of empowering in the model. The model focuses on empowering scientists with the resources and flexibility to pursue breakthroughs. That starts with our programme directors. As you know, this is a group of scientists and engineers who we have recruited from across the UK and beyond to shape the funding programmes. The job of the programme directors is, first, to identify areas where scientific breakthroughs that can deliver a step change in economic and social prosperity are unlikely to occur. There is something very consequential, but it is unlikely to occur because the R&D that is needed is too speculative, bold, maybe interdisciplinary, or the current system is not set up to combine the elements that are needed to deliver that breakthrough. So our programme directors define those areas, and then their job is to connect and fund researchers across the UK and beyond who can seize that opportunity.

Baroness Willis of Summertown: Who then kicks the tyres? Who checks and says, “This is a novel, exciting area”? Where is that scrutiny?

Ilan Gur: When our programme directors started at the agency, it would have been easy to say that their job was to define the



HOUSE OF LORDS

funding programme. We said that, actually, that is not the job yet, because everything we are doing at ARIA will be quite speculative. We are going to be taking a first step, and we want to make sure that the step is not taken in an unproductive direction. The first job of our programme directors was to define what we call “opportunity spaces”.

Each of these is a critically important but underexplored area of research that we believe is ripe for breakthroughs, and our criteria are that there is something highly consequential for society and the future of the UK here. Relative to that importance, this area is underexplored, and we think it is ripe for transformation. Refining these areas through a rigorous discovery process, we are talking to the research community and publishing openly the thesis of these areas, getting feedback and refining and evolving them. We drive from there to a broad area and then to a research programme that is very specific and has a very clear target.

Through those programmes, we have now launched seven of those. Those are funding programmes aimed at answering questions such as: can we demonstrate AI hardware at 1,000th the cost and energy consumption of today’s state of the art? Can we leverage the UK’s strengths in synthetic biology to completely reimagine the global agriculture industry and how agriculture can serve human and planetary health? Or can we combine our expertise in the UK with, for example, earth scientists now modelling AI to forecast and prevent climate tipping points?

We get feedback for the thesis for the programmes we publish, and before we allocate any funding to the programmes we go through a rigorous internal review and approval process for the thesis of the programme to make sure it aligns with ARIA's mission and mandate and that it is been prepared responsibly.

Q4 **Baroness Willis of Summertown:** I am a professor of biodiversity science, so a plant scientist by background. There are some areas in your plants programme—I cannot comment on it because it is not my specialism—like the C4 rice project, which have been extensively funded by the Gates foundation already. So how does that align with what you are going to be doing?

Ilan Gur: That was a wonderful question. When we approve a programme internally, we have four criteria for that approval. First, is the programme worth shooting for? The idea there is that, if we were to succeed at meeting the objectives, do we think there is a reasonable chance that this can be absolutely transformative?



HOUSE OF LORDS

Secondly, do we believe the programme is differentiated? That is where we dig very strongly into how this compares with what is being funded in the system today and how it is being funded. We really need to make sure that we do not believe there is overlap and redundancy; I can share an example on the plant programme related to that.

Thirdly, it needs to be well defined and clear. If we are going to put out a solicitation for funding, we need the research community to understand that this cannot be arbitrary or change every so often. We need to be very clear up front.

Fourthly, has it been responsibly designed from the standpoint of regulation, governance and social ethical responsibility?

Angie Burnett, the plant programme director, who is a plant physiologist by background, went through a number of iterations of what that programme could look like. Ultimately, she ended up with a programme the thesis of which is that the UK is extremely at the cutting edge of synthetic biology. So much more funding and attention has been paid to synthetic biology for mammalian systems, yet the power of plants as a technology platform for humanity—I do not have to tell you—is unbelievable.

What differentiates that programme is the boldness of its target. It aims to create, for the first time, a synthetic organelle within a plant system that can be maintained and transferred within species. There is related research but, by and large, the community believes that, if anything, this is really out there, and we are very excited about what it might mean if it works.

Q5 The Chair: You have talked about reducing the compute energy for AI and the challenge of reducing it by a factor of 1,000x, which is great. Following up on what is failure, what might make you stop something if it looked as though it was only going to get to 100x? Would you keep on funding it, or would you say that it was not getting far enough for you?

Ilan Gur: That is a wonderful question. The programme on unlocking AI compute and scalable AI compute is run by Suraj Bramhavar, one of our programme directors, whose background is as an electrical engineer in research, and he had done a start-up. In formulating the programme, he chose this goal of a 1,000x improvement, because it was clear, when he talked to global hyperscalers and AI companies, that that would be enough of an improvement to get the global industry to think about shifting some of the core assumptions about the supply chain and where investments might go. It was also clear that it was so far beyond



HOUSE OF LORDS

the edge of what they are doing now that it is unlikely to happen anyway. We are now funding, through ARIA, a number of projects that are working in co-ordination, and competition in some ways, to try to achieve that goal. Ultimately, Suraj is now directing that programme in close co-ordination with those researchers.

If we find out that a project is not going to hit 1,000x, but only 100x, should we continue to fund it or not? The idea will be to go back to those first principles. Do we still believe that, unless it is a 1,000x, it is not going to change the conversation about what is investable or valuable to the corporates in the world? In that case, we will probably discontinue that project if we think there is no way it will hit it, or we will redirect it towards a new approach. Sometimes by trying and not quite reaching it, you will open up a new direction, but ultimately it will be Suraj's job, as programme director, to make those decisions along the way, and likewise the other programme directors.

Matt Clifford: This is a good opportunity to talk about the stage at which ARIA comes into the equation. Clearly, if that 1,000x improvement is achieved, which would be an enormous transformational accomplishment, it will require a lot of capital that ARIA does not have. In many ways, we probably will not know whether it is 10x or 100x or 1,000x improvement during the period when ARIA is funded. We will know in the case of clear dead ends but, in many ways, the question will be: can we find the right scale-up funders to come in after us? As you know, ARIA is designed to not take things all the way. In some ways, the point of all the community building and engagement with stakeholders during the programme formation is to ensure that we have at least some plan for when ARIA stops funding it, whether it is because it has reached the end of that stage or because it is not working. Will there be someone there who will have the incentive to pick that up?

Q6 **Baroness Young of Old Scone:** Baroness Willis stole some of the questions I wanted to ask. I wanted to ask about the model, your researchers and programme directors. You have talked about how your programme directors define their missions. Tell us a bit about how you work with researchers and whether you have found new ways in which to incentivise or fund them. What have you learned from the experience so far of funding people for projects that are on the edge? You talked about investing in people rather than in projects. How do you find those researchers if they are researching an area that has so far been underresourced or not resourced at all?

Ilan Gur: I will address the second of those first, and then we can loop



HOUSE OF LORDS

back to what we have learned in the funding. One of the things that has been exciting to see is the way the opportunity spaces and programmes our directors have shaped are attracting talent in new ways in which to go after those problems. It is worth thinking, from a first-principles perspective, about how we normally fund R&D from government for society. Oftentimes there is the approach: “We have a well understood problem that we think is important to our constituency, to society, and we’re going to look for research that can solve that problem”. You can think about that as sort of a top-down approach. Another is to say, “We have brilliant scientists, and we know that by providing them with freedom and space to do their work, we can’t prescribe what’s going to come out, but incredible innovations will come out of that. So let a thousand flowers bloom”. Those modes of funding exist and are the basis for a lot of the powerful base of research in the UK.

ARIA is meant to be doing something different and create a model that can change some of the outcomes. One of the ways I like to think about this is that, if you want to do things differently, you need people who see things differently; you need a lens through which to look at the problems differently. Our programme directors, when they shape one of these opportunity spaces or theses, are taking their unique experience informed by lots of conversations and are developing a thesis that looks different from what exists in the system. It is not simply, “We’re going to fund a discipline such as biology, or a problem”. It connects the two in interesting ways.

For example, Jacques Carolan, one of our programme directors, was a quantum scientist who used to build quantum computers. On coming into ARIA, he decided to shape a programme and work in neurotechnology and neuroscience because he felt that it was an incredible opportunity; if we think about the broad range of neurological disorders—Alzheimer’s or Parkinson’s, ranging to mood disorders and depression—the burden to society is enormous in the UK.

We have learned scientifically that the broad range of these disorders relates generally not to a specific part of the brain but to disorders of neural circuits. Different regions of the brain take part. There are different cell types and time constants. Imagine a technology platform that could interface, understand what is happening at that circuit level and potentially intervene. There is the potential to develop a single technology insertion point that could address this broad range of disorders, which, in an integrated way, represent about three times the burden of heart disease. Think about the expected value of such an invention or technology suite. It is a pretty significant portion of all of global pharma.



HOUSE OF LORDS

Jacques realised that that circuit-level brain work is not really happening. Why? Because this is not just about neuroscience but about clinical surgeons, devices, electronics and bioengineering, disciplines that are hard to cross paths in. It is also a question of academia, as well as start-ups and hospitals. So when you look at the space that he has defined, elements are highlighted where these different communities see themselves in the problem in ways they had not done before.

This is a long-winded answer but one that hopefully brings to life the idea that in putting that out there we are attracting researchers who we never expected, who say, "Wait. I think I might be able to work on neuroscience and neurotech in ways that I didn't expect". A great example is the cutting-edge AI that we think is needed and where there is an opportunity to understand the brain-state dynamics within these regions and neural circuits. This is unexplored territory, and there is a chance to do that in ways that no one in the world has done before. That programme is now in the process of negotiating awards to exactly that—everything from bioengineers to AI scientists to biologists across these different parts of the ecosystem.

Baroness Young of Old Scone: Where does the join-up come? Do the researchers assemble these cross-disciplinary groups themselves, or are you funding siloed research areas which you then amalgamate together?

Ilan Gur: It is a great question. We really want to be driven. We talk a lot about leveraging the intrinsic motivation. It is about empowering the scientists. We see ARIA's role as being a catalyst, an intermediary. We found in the programmes that by publishing the thesis we have people sharing those documents with others in their community and outside. We then, for each of our programmes in development, held a workshop. I will share another example here, which I love. Angie Burnett, the director of the programme on programmable plants that I mentioned, held a workshop that brought together synthetic biologists and plant scientists across different institutional types. It was held in Bristol, I believe, and we had about 65 people in a room, including some international experts from outside the UK.

As part of the introduction, I asked the audience, "How many other people do you know in the room by show of hands? Do you know more than half?" The vast majority knew fewer than five people in the room before they showed up. Yet the buzz and the collaboration were so vibrant that it was a signal. We are creating and catalysing some new connections. We found that the



HOUSE OF LORDS

applications that have come into our programme include collaborations from cutting-edge synthetic biology labs that a year ago would have said, "There's no way I'm working on plants. It's just not worth it or interesting. There's not enough funding". Now they say, "Wow, this may be the most important thing I do", and vice versa.

Q7 Lord Drayson: You described applications that could have significant economic potential if they come off. How is the intellectual property handled? Part of the rationale for Parliament agreeing to set up the agency was that it would provide strategic advantage to the United Kingdom. For example, you are funding some work at Cornell on backpropagation-free algorithms. What rights does Cornell have to anything that comes out of that?

Ilan Gur: This is a great question. We have tried to apply to every part of what we have set up with ARIA the idea that we should be building not based on precedent but based on principle. The ARIA Act is where we start: what was the mandate, and how should our policies, our contract terms and the way we develop programmes flow down from being able to accomplish what we were set here to do? At the same time, we have this rich history of agencies that have operated differently and transformed the world. Obviously, DARPA is a great one.

Lord Drayson: Given the time constraints we have, could you just answer the question?

Ilan Gur: Yes. We spend a lot of time thinking about contract terms and IP in particular. I will share a few elements of that. One is that we do not want to be seen as competitive with any other part of the system, because our job is to be an intermediary and we want the research we fund to lead to great impact. So, as a default, the research institutions that we fund keep the IP.

Lord Drayson: Right, so it stays with Cornell in that case.

Ilan Gur: A big part of our mandate is the need to show clear benefits to the UK. You want to drive to scale through allowing the market to operate, so how do you make sure that the UK gets benefits? Through a lot of early engagement, there was a sense that if you can increase the probability or level of impact of one of these programmes by 10x or 100x by funding some portion of the programme outside the UK, because there is a unique capability there, we should not hesitate to do that.

So we are doing that with some programmes, but our IP terms are such that, if the IP ultimately ends up being licensed to a non-UK entity, there is an additional royalty payment stream that will come



HOUSE OF LORDS

back to the UK. That is one of the ways we have tried to address that.

Lord Drayson: Has that been published?

Ilan Gur: Yes. Similar to the programmes, we published the contract terms, got feedback from many institutions around the UK, and have updated since then.

Lord Drayson: So the IP is handled in the same way, whether it is a UK sovereign-based researcher or someone outside the UK. The IP is not retained by ARIA.

Ilan Gur: The IP is not retained by ARIA. That is very important. Ultimately, we see that as a way to drive the market forces to get the maximum impact and to make sure that it is not ARIA that succeeds here but the UK research ecosystems that are developed in these spaces.

Matt Clifford: To be clear, the royalty fee is payable only if it is a non-UK entity, not if it is a UK entity.

Q8 **Baroness Willis of Summertown:** I want to get a tiny bit more on the specifics of the funding model and how it compares to the UKRI funding. Some of the examples are fantastic, but many of them are relying on universities to deliver that science. Do you have the same funding model with an 80% overhead that goes to the universities? How does it work?

Ilan Gur: The specific answer here is that we fund 100% of the research in our grants². There are a few reasons for that. One is that we know that the programmes we are funding are very different and at the edge of what is possible. They may fail, but we may pivot them. So there is a different burden on the institution. Again, we have our own IP terms that look differently, but we are also driving things with quite a bit of urgency. We wanted to make sure that we were not requiring, say, match funding, which would suggest that another institution saw this as a priority for them. This sort of ensures that what we are doing is differentiated, and that we can move at pace.

Baroness Willis of Summertown: So you fund 100% of overheads, do you?

Ilan Gur: Correct³.

² Note by the witness: Ilan Gur was confirming that ARIA funds 100% of project costs (not necessarily a 100% overhead recovery rate - as overhead recovery rates may vary by organisation)

³ As above



HOUSE OF LORDS

Q9 Baroness Northover: These questions are about a series of things: your theory of innovation, among other things, which I look forward to hearing the answers to; blue-skies research versus applied research; and the level of invention.

First, are you really doing blue-skies research, or is there a clear pathway to some application or invention at the end of it, and what do those pathways look like? You have answered that in some ways already, but could you comment further on that?

Ilan Gur: I would not describe what we are doing as blue-skies research. Most importantly, our funding is driven towards creating some new capability. One of the differentiated aspects of ARIA is our ability to direct these programmes and learn and adapt. If this was just creativity-driven blue-skies research, why would you need ARIA to adapt that? You would just let the researchers do it themselves.

Maybe the most helpful thing here is the phrase we use for a successful programme: that by the end of a programme, if it is successful it would have changed the conversation globally about what is possible or valuable in a space. That change in the conversation means that we have shifted the trajectory of what might be possible towards a capability that we believe has a reasonable chance of being enormously transformative for society. Hopefully that helps.

Q10 Baroness Northover: Yes. To dig a bit deeper into the theory of innovation, the summary here is that, if you have clever people with money and the freedom to pursue a mission, they are going to innovate. The question is whether innovation can work like this, given increasing scientific sophistication, specialisation and the need for massive collaboration. In some ways, you have answered that question, because you are seeking all those things within innovation, but, again, could you address how you think innovation occurs?

Ilan Gur: We have a lot of examples, for instance from the ARPA agencies or DARPA in the US. One of the biggest things for me that feels universal—to Matt's point—is to make sure that the agency is always defining very bold challenges that are outside the norm of what is being funded, because as long as you are uncovering something at that edge you have a chance again to change the conversation about what is possible.

Ultimately, for us, one of the core principles behind the theory of change is that these technology breakthroughs on their own do not deliver the mission that you have set for ARIA. A technology breakthrough might change the perspective of what is possible, but



HOUSE OF LORDS

we succeed only if we deliver lasting value for the UK in profound ways.

If you work backwards from that, what will success look like? It will be entire new technology platforms that improve people's quality of life, entire new industries that are born in the UK. For that to happen, there needs to be a point at which this changes the economic conversation: where is investment flowing? What are start-ups or big companies doing? In order for that to happen, you need entire new communities of practice and talent to flow. We see this as a chain from technology breakthroughs rolling into changing the conversation in terms of talent and economics. To the member's comment, the idea and the intent behind all these programmes is that, if they are successful, all that will be happening in a way that is rooted in a new UK ecosystem.

Matt Clifford: I just emphasise again that ARIA's role here is very much catalytic. I do not see its approach as being at odds with the idea, which I agree with, that a lot of this stuff requires large-scale collaboration. One metaphor we use quite a lot at ARIA is that an ARIA programme shoots up a flare about what is possible. That flare is designed to attract a very broad, diverse, even international ecosystem of collaborators. So I would not see it as a theory of innovation that says that it is a lone ranger that does it by itself. It is this idea of community, of catalytic activation energy. That is at the heart of everything we are trying to do.

Q11 **Baroness Northover:** Okay, so this is about one that works out, as opposed to one of the ones that has failed. You have addressed this to some degree, but how do you spin them out into larger research institutions, and how do the activation partnerships work in that regard, or how do you hope they will?

Ilan Gur: One of the things we have tried to be very deliberate about is to think proactively and pre-emptively about what the pathway will be when we uncover technology breakthroughs is what the pathway will be through which they end up moving forward.

We expect that, for some programmes, the pathway will be that that new capability is very strategically important or compelling to a part of government or to large companies. However, across the board for many of these spaces, we see science entrepreneurship as one of the key drivers for moving things forward.

Activation partnerships is an initiative we have launched that is not thinking about creating specific pathways pre-emptively, because it is too early—these are all quite speculative—but is figuring out



HOUSE OF LORDS

ways to inject and create the conditions within our spaces so that, when technology breakthroughs emerge, they have the maximum chance of moving forward, because commercial mindsets, insights and entrepreneurial talent are really plugged into that space.

Baroness Northover: You are bringing those people in to be involved as part of that collaboration.

Ilan Gur: Yes, and I will give you one example. Activation partnerships, like many of the things that we are doing, is an experiment. This is the first time we are doing it, and we will learn. Some of those activation partners are entrepreneurial venture firms thinking about new talent programmes here in the UK for the first time across several spaces. One example, Cambridge University Health Partners, applied to become an activation partner and work with us to think about how we take the work that is happening on the precision neurotechnologies that I described, these new neural-circuit technologies, and make them lead to something that touches someone in a clinical setting.

There are enormous questions there to do with regulation and how to navigate it, patient engagement, and the institutions, hospital systems and economic frameworks. They have assembled groups from around the country to come together not just to think about making ARIA successful but to think alongside us. The phrase we use is that they will be “running alongside us” in this space, thinking about all those other aspects and plugging and tying them into the technology, which is very exciting.

Q12 **Lord Rees of Ludlow:** I would like more clarity on who will do the work and employ them. I guess that in most these areas—plant science, neurology, et cetera—most of the people are in universities and research institutes, and will contribute as part of their own team perhaps. It is not completely clear to me to what extent they will be full-time employees of ARIA for a period, how that will work and where they will be based. Could you clarify that a little?

Ilan Gur: Thank you for the question. ARIA’s role is as a catalyst, an R&D funding agency. We are not doing any of the research ourselves organisationally. Our programme directors are thinking about the research to guide it, but our research-granting contracts are entirely arm’s length and go to other institutions. As you said, in many cases those are universities. We have contracted maybe 32 universities for research so far, but each of these programmes fund what I think of as a constellation of projects. In almost all cases, the universities work alongside large companies, small businesses or start-ups, and, in some cases, individuals or entirely



HOUSE OF LORDS

new organisations. We are agnostic to the institutional environment and, in fact, want there to be a diverse institutional network.

Lord Rees of Ludlow: From the perspective of the people doing the work, they are on a grant but, in effect, it may be a more permissive grant than they would get from a research council or university. Is that the way to look at it?

Ilan Gur: Precisely. They are on a grant or a research contract that, in many ways, looks different from the ones that they might get from another research council.

Lord Rees of Ludlow: Will it be channelled through a university, or will they take leave or secondment from their university while they are on this grant?

Ilan Gur: The default is that you apply from the institution from where you will do the research. That is where you are employed and will do the work. An exciting thing we have found is that, so far, our grants have catalysed some mobility from researchers, who are saying, "With this grant, I feel like the right environment for this research might be outside the institution I've been working in", but that is for a minority of cases.

Q13 **Baroness Neville-Jones:** You mentioned DARPA a moment ago, and a lot of the comments on the origins of the thinking about ARIA liken it to DARPA. Obviously one of the things about DARPA is that it has a big master, with a big budget, which can do the pull-through. There is no such relationship for ARIA. How do you create a situation in which you can get a GPS system out of the technology and the science that lie behind it? It seems to me that it is still a UK problem about scale-up and getting people to work together. Do you have a way forward on this, which is a secret that we have not hitherto managed to crack?

Ilan Gur: One of the big differences between our mandate and DARPA's is that we are not focused on defence and do not have the Department of Defense procurement budget driving things.

I have experience from being one of the founding programme directors of an agency called ARPA-E in the US. That was also inspired by the DARPA model but focused on climate and energy. That was an example of creating an agency like this, which is focused on sectors where the Government would not be the customer and you would need to leverage the private sector.

One of the key lessons from ARPA-E, which has been successful in a number of areas in catalysing innovation in climate and energy, was the power of entrepreneurship as a driving force. In some



HOUSE OF LORDS

ways, we have found that ARIA has an advantage here, not just because of the strength of the UK talent base in the ability to activate people to be more entrepreneurial in their work, but because our mandate is focused on not only defence as an application. We are seeing talent that is wildly enthusiastic about pursuing these areas, because they are motivated to figure out how to improve life in the UK, for society.

Baroness Neville-Jones: Where are you finding that talent?

Ilan Gur: Across the board. They are coming out of universities at the post-doc level all the way through to senior professors. There are start-ups and entrepreneurs here in the UK, but also from outside the UK. People are attracted to contributing to a programme and spending a lot of time and work in collaboration with UK entities. In some cases, we have new organisations and people creating their first offices here in the UK through these programmes. It has been an exciting start.

Matt Clifford: The list of grantees from the scaling compute programme, which has been referenced a couple of times and was published two weeks ago, is really quite inspiring. You see that the breadth of organisations that are receiving grants in this space includes a number of start-ups in the UK and at least one non-UK organisation that is building a new subsidiary in the UK in order to base its research here.

This idea that scientific entrepreneurship needs to be at the heart of our scale-up mechanism is quite radical and new. If it works, it will provide a really exciting model for thinking about how we play to the UK's strengths, given that we do not have \$900 billion a year.

Baroness Neville-Jones: Are you in a position to give the Government a kick, when you feel that something can be procured and really developed? Our problem is a lack of take-up, and government is as guilty as anyone. Our own Ministry of Defence is not known for backing UK invention or development, so do you envisage creating a climate to change this reluctance to come forward with some public money?

Matt Clifford: I have two points on that. First, at a broad level, engagement with ARIA across the public sector has been fantastic. If that was the right path, we would have no shortage of connectivity into the system to do that.

Part of the design of ARIA, of course, is to have the Government Chief Scientific Adviser on the board, and that has already been



HOUSE OF LORDS

really helpful in that connectivity. I would emphasise that there is a scale issue in terms of government as customer. That is not the core of these programmes. In fact, the whole purpose of scientific entrepreneurship at scale is to ask what problems in the global markets need to be solved. If you look at some of the examples Ilan has used, such as neuro-precision tooling, this is a global market, and we need to find that global pool. To your point about scale-ups, I have spent my entire career in that space. You are right: that is a big gap. But we need to think globally about the customers for these things, and in the reach of the board, the programme directors and the team we are well equipped to do that.

Q14 Baroness Neville-Jones: I would not dream of arguing that there is no global opportunity in the ability and the willingness to procure, but there are weaknesses in our system of a kind that are well known. How well placed are you as an institution to get parts of the system that have been quite backward in coming forward to take advantage of what you are promoting?

Ilan Gur: I have a few perspectives here. I believe that the kick to the system will come from the technology breakthroughs and proof points that emerge from these programmes, and from the attention that they are getting globally. In some areas, we will see in the coming years that the potential value and the strategic differentiated advantage of the UK is profound. From an optimistic perspective, one of the benefits of ARIA is to be able to build with this decadal mindset. My expectation is that a number of barriers that have historically prevented maybe entrepreneurship or procurement or scale-up will, five to 10 years from now, through the work of many other folks across the systems, get better and the conditions will be ripe. In some cases, we are trying, like our activation partners, to help in our spaces to get the right conditions. There are so many great actors working towards that.

Baroness Neville-Jones: Is the private sector and private industry coming through?

Ilan Gur: One point of evidence here is that, through this activation partner call, we have two world-class, global, seed-stage venture capital firms that are starting activities in teams in the UK for the first time. We are not funding and working with those firms; we are working with their new UK entities, and another global venture group likewise. There is a sense of a moment now where the opportunity in the UK is ripe, and we see ARIA as one of many things that are catalysing that, which is great.

The Chair: We hope it is as positive as that.



HOUSE OF LORDS

Q15 Baroness Neuberger: I need to declare an interest first, because of your neuroscience area. I chair University College London Hospitals NHS Foundation Trust, which includes Queen Square and the National Hospital for Neurology and Neuroscience. I suspect you may have been talking to some of those people, which is why I ought to declare the interest.

I will pick up on Lady Neville-Jones' question in a moment, but I want to ask you about your enormous ambitions and missions, when you do not, in the scale of things, have a very large budget. Do you have any worry that the budget is just too small to make the difference? Or is what you are saying about the venture capital people coming in the answer to this? In other words, it is the outside players who will make this possible.

Matt Clifford: I go back to this idea of ARIA as a catalyst. If you think of the funding journey that any of these programmes will go on, clearly only a fraction of that funding will come from ARIA. We hope we demonstrate, in the coming months and years, that ARIA is modular and therefore scalable. Today, we are running seven programmes and have allocated our budget accordingly. We obviously believe that we will have enough success across those—as I said, even if it is one. In fact, I would be surprised if it was more than one, given the risk that we are trying to take. I do not think—hostage to fortune—that we will ever say that ARIA would like to invest 100 times as much in a single programme. That is for other people to do. But we can certainly imagine a world where we have more programmes, because the programme director is the unit of scale. We are in the middle of recruiting our next cohort of programme directors, and we can say very clearly that there is no shortage of talented people we could appoint to be programme directors.

Q16 Baroness Neuberger: That is very reassuring. Why do you think DARPA has been such a success? How much of that will you emulate, given that the budget is so much smaller? How will you differ? And are there other innovation agencies you would emulate, given that you are not on the scale of DARPA?

Ilan Gur: The real thing we are trying to emulate from DARPA is the idea of cutting through the way the system normally funds by having programme directors who can connect the dots in new ways. That also suggests a bit of the answer to your last question, which is that by connecting the dots in new ways we are not funding in a vacuum and trying to go from zero to full success; we are activating the system that exists. ARIA succeeds only because there is such a rich base to build upon. That is how we should imagine that ARIA, hopefully, can have an outsized impact. Even if



HOUSE OF LORDS

the budget might, in absolute terms, look small to certain people, what does it look like in relative terms, and in what ways does it leverage the rest of the system?

Q17 **Baroness Neuberger:** Fair enough. How do you communicate with wider society? I want to include in that why freedom of information legislation does not apply. That is also important.

Ilan Gur: Why freedom of information does not apply might be a question for the parliamentarians who made that decision. My interpretation of it is that ARIA really needs these programme directors and our team to be focused on the search for the breakthroughs that can be consequential in as efficient, effective, dynamic and agile a way as possible.

We have built ARIA in a way that suggests that, really, we want to be building in public. That surprised many people, but our view is that the advantage, because we in ARIA are not focused on defence or anything secret, is that we have the opportunity to publish our theses, get incredible perspectives and engage talent that would not otherwise be engaged. We have found that to be valuable and productive. I believe we will continue to bias towards being very open in what we are doing, who we are funding and what we are learning along the way, including where we might have to shut down projects or pivot based on learnings.

Baroness Neuberger: So it may be odd about FOI, but nevertheless you want to be open.

Q18 **Viscount Stansgate:** When Parliament was debating the setting-up of ARIA, it was in effect agreed that you would have a licence to fail, which is a relatively new concept. We have heard from other witnesses about the high risk and reward you were hoping to achieve. You have partly answered my question already, but can you just very briefly outline what you think success for you looks like and how you would measure it?

Matt Clifford: One thing I think about a lot is in what ways the private sector is good at this and which elements of that we are trying to replicate. My professional background is as a venture capitalist, and the way it works there is that you expect that most of your investments will fail, but the ones that succeed will pay back many times over. That is harder to achieve in the public sector, particularly in a model like ARIA, where we are not taking shares or equity in the things that we are funding, because in venture capital you can count how much money you are getting back and how much you put in, and that is how you know whether you are succeeding.



HOUSE OF LORDS

So although we do not measure the impact in the same way, through cash returns, I think we have the same mindset, which is how we think about the total amount of funding and resource that ARIA has spent and, over the long run, how we look at the total impact that it has had.

I go back to the idea—as a board, we push the execs on this all the time, and I think they have done an extraordinary job—of saying that we will only fund programmes that, if they succeed, would justify ARIA’s existence if they were the only thing that succeeded, measured in economic benefits, health benefits, quality-of-life benefits for the UK. There are many ways in which you could measure that, but that is the way we have defined what success would look like.

I like the framing that Ilan used earlier about changing the conversation, but it is really saying that, if this worked and was the only thing that worked, would you, the public, the media say, “Yes, it was worth having ARIA simply for that outcome”. In the example that was raised earlier of DARPA, you could argue that the creation of the internet would have made DARPA worth it even if nothing else had happened. That is the aspiration for what good looks like at ARIA.

Viscount Stansgate: Do you have any sense yet of whether you are helping to promote something that could lead to such a transformation?

Matt Clifford: I certainly feel that each of the seven programmes we have, if they achieved their goals, would pass that test.

Q19 **Lord Wei:** I also declare an interest, being involved in venture capital through groups such as Future Planet Capital. Many of our witnesses were very supportive of ARIA but wanted other parts of the research ecosystem to fund high-risk, high-reward projects. What do you think UKRI, or parts of it, could learn from how ARIA has worked?

Ilan Gur: We should recognise that this language of high risk and high reward is flawed, in the sense that all pursuit of science is high risk and, we believe, high reward. UKRI is certainly doing high-risk research. The key is that ARIA is set up to do it in a different way and in a different mode. Our programme directors, our team and I have benefited from interacting closely with UKRI, the other funding bodies and learned institutions to share what we are learning about the unique mode that ARIA is introducing into the system and how that can inspire or inform the ways in which UKRI does things, or vice versa.



HOUSE OF LORDS

The image to have in mind here is one of symbiosis. If, over time, ARIA starts to look more and more like other funding bodies, that is a problem. If we start to integrate into something that looks the same, we will be losing the symbiosis, the hybridisation, the new modes in the system and that diversification. That is the way that we see it.

Lord Wei: You spoke earlier about the approach that you take and how some of your programmes are incredibly ambitious. You also talked about your risk appetite. Can you tell us more about how you assess the difference between something that has a 5% chance of success but could produce an amazing breakthrough, versus other projects that have a higher probability of being successful but may not be such a breakthrough? How do you make that judgment?

Ilan Gur: It is a great question. There is some humility that comes with this work, which recognises that the error bars on our predictions or what we believe might come from this are very high. As Matt said, we want to ensure that these are worth shooting for, such that, if we succeeded, they would be incredibly important.

Whether we think we can succeed is an interesting question. One of the things that I like to share might be helpful. When we do engagement to develop our programmes, I tell our programme directors, "I should be able to poll a set of experts, once we launch this programme, to ask them what they think. If they all think it's brilliant, we've done something wrong, because it's likely to be happening anyway. If they all think it is ill conceived and will never work, we've done something wrong, because it's probably just a bad idea. But if I get a variance—if half of them think it's brilliant and half of them think it's a total waste—we may be on to something, so it's worth taking a shot".

Q20 **Lord Wei:** Let us talk about failure. Look at the US space programme, for example, or even some of the work that DARPA does. Sometimes good things can happen from failing to reach objectives. How do you look at failure constructively? What will you measure when programmes "fail" to capture any of the benefits that might have happened none the less?

Ilan Gur: You have said it. As long as the pursuit drives learning and adaptation of our assumptions and strategy, we are succeeding. Very importantly, each of these projects is not happening in isolation. We have a programme director who is capturing this learning and guiding the next steps. With many of these programmes, it is worth mentioning that we may get to the end of a programme and it may have failed to meet its objectives, but it may have inspired the next programme director to say,



HOUSE OF LORDS

“Here’s v2 of the programme”. We would have learned so much and seen an incredible galvanisation of talent and institutions that want to make this happen. That is our approach.

Q21 Lord Lucas: You have talked about the relationship with other research funding sources. How would you characterise it? Are you getting collaboration or genuine, enthusiastic partnership? What has been the best and the worst?

Ilan Gur: I will have to think about the worst along the way. It may be helpful, because we have talked about this a bit, to just give an example. We have a programme on forecasting climate tipping points. One interesting thing about this programme is that it is being led by two programme directors, Sarah and Gemma, who are on secondment as university professors, but their university research has focused on human health. They are medical physicists who build devices for diagnosis in the human body. They came to ARIA and said, “We think climate change is critical. As much as there’s talk about tipping points, it’s very clear that the race for decarbonisation is a race against time for the consequences of climate change”.

The innovation community is not really thinking about tipping points. Do we even understand them and know when they might happen? They developed a programme in this area that allows them to bring in new research and insights. They found an opportunity to join the optics community with the climate community in many new ways. As they shaped the programme, part of their learning was to go to UKRI and the research councils. They were learning from NERC. But, in shaping the programme, they realised that some parts of it needed infrastructure that already existed and that it would be redundant for us to create it anew. That programme has formal partnerships with the Met Office and the National Physical Laboratory, which are bringing their resources in a catalytic way to support all the other researchers we are funding. That is an example that we hope will go well.

In terms of what has perhaps not gone well, this is more about the ecosystem as a whole, but we mentioned that a number of our contract terms are different in relation to the UK benefit and other aspects. There was friction in getting institutions to become familiar with those for the first time. There were a lot of questions and we had to spend a lot of time explaining the rationale. That created a bit of an activation barrier, but it has also given us the benefit of saying that we now have 32 universities that have signed up to our contract terms and world-leading entrepreneurs, venture capital firms and so on. Part of this is just a learning and adapting process.



HOUSE OF LORDS

Q22 Lord Berkeley: I am interested in the projects that might seem a bit more controversial. Is there an ethos to do the bits of research that might be hard to fund elsewhere? Some of them might well need ethical or regulatory oversight. How do you bring that into the equation? Does it work, and are there examples of where it has gone wrong or could go wrong?

Ilan Gur: We think about this a lot. When you are working with frontier technologies and research at the edge of the possible, it brings you into territories where it is incumbent on us to think about not only what would be the benefits but what might be the negative consequences that could come from the research.

We have a layered approach to trying to make sure that we are responsible in our governance. The first I have already mentioned, which is building in public. Shedding light to the public on everything that we are doing allows us to get a clear signal of where the concerns are and what their amplitude is from different stakeholder groups. Then we can be deliberate with our engagement to engage with those groups, to learn and to adapt. We have seen that, for example, in our cooling the planet programme, climate engineering. We had several adaptations of the thesis based on non-technical input and the responsible governance of that.

In terms of the review of projects, we have criteria-based expert review with a number of voices. Finally, we are engaging in certain programmes with social and responsible governance and ethics experts, either as advisory boards or to oversee the projects that are working.

Lord Berkeley: So you are not worried about public backlash, and you have some ideas about how to deal with it if you get one, because you might have to close a project down if it got really bad.

Ilan Gur: Are we worried about public backlash? We think about that and about maintaining trust. When we are designing these programmes, the real question comes down to our principles. Do we believe that the knowledge we could create through this programme is enormously important and consequential? If so, how do we take as many voices as possible, with input and lessons from the past, to take a shot at uncovering that knowledge in the best and most responsible way? That is the best we can do, and then we will learn and adapt.

Matt Clifford: One very important decision that we made about the way in which governance would work at ARIA is that the board would not directly intervene on individual programme-level



HOUSE OF LORDS

decisions. The way we see it is that strong oversight of the executive leadership is strong oversight of the programmes. In fact, one of the key things about getting governance right in an organisation like this is providing the right principles and checking regularly on those principles but not getting into the weeds of individual programme decisions. We have established an ethics subcommittee that is able to provide continuing oversight on programmes, particularly those with potential controversy and ethical complexity. The role there is partly to ensure that we are not taken by surprise by any of the sorts of things that you are talking about.

The most obviously controversial programme so far is Mark's programme, which Ilan mentioned, on climate interventions. My view, and we have discussed this many times at board meetings, is that this has been handled in an exemplary way by the executive team in being very open and transparent, engaging stakeholders from across that ecosystem early, and adapting as necessary in response to feedback and providing regular updates. We mitigate the risk in that way. It is important not to let fear of hypothetical backlash constrain the ambition but to monitor that and move incrementally in response to what we hear.

Lord Berkeley: And to have a plan for dealing with it if it looks as if it is going wrong.

Matt Clifford: Absolutely.

Q23 **Lord Drayson:** The committee is very supportive of ARIA's mission and wants you to succeed. My question is about securing your future. I think I am right that your Treasury funding goes through to 2026, but your parliamentary review is on a 10-year basis, in 2033. So how on earth are you going to be able to secure funding beyond 2026?

Matt Clifford: Clearly, the upcoming multiyear spending review will be a really important moment for ARIA. We are starting from a position of strength in having broad cross-party support for ARIA, right from its inception. It has been great to see with the new Government that the depth of support in DSIT and beyond for ARIA is extremely strong. The ARIA mission has been warmly embraced by our new Ministers. It probably helps that our new Science Minister was one of the founding board members of ARIA, for example. You are absolutely right: this is one of the key challenges and the support of stakeholders from within Parliament, the Government and the ecosystem is going to be key.



HOUSE OF LORDS

We recognise that ARIA, over the long run, has to provide great value for money, just as any use of public funds needs to, but that will obviously be measured in a different way, with a different risk appetite, and over a different time frame from the way many other uses of public funds would be. That is why we are so keen to establish now these core ideas about what failure and success mean for ARIA, what proportion of our programmes we expect to succeed, because we do not want to end up with a situation in a couple of years when people say, "An ARIA programme failed. Does this mean that ARIA is failing?" We have a strong narrative with broad buy-in.

Q24 Lord Drayson: That is good to hear but, to push you a bit further, the present Government's number-one priority is to deliver economic growth. As you have already heard from other members of the committee, it has heard multiple witnesses in many inquiries telling the sad, long story of UK science making breakthrough contributions to innovation but not being exploited and creating economic growth here. You basically invoked the power-law argument to your strategy. How do you make the argument to the Treasury, taking a long-term view, that that will not happen with regard to ARIA-funded projects, that they will lead to a movement in the dial of UK economic growth because one of them has made breakthrough developments that have led to a strategic advantage for the UK and economic growth in the UK?

Ilan Gur: Matt and I, in different settings, have been involved in early-stage innovation projects that have led to multi-billion unicorn outcomes at a company level, based on cutting-edge science—and, in my case, in entirely new industries, if I think about some of the work we did at ARPA-E. We can look at the expected value, economic and social, that can come out of the programmes we have already launched and the ones that we anticipate launching in the next year, add that up and look at the early signs—the global talent that is being attracted to these programmes, and new ways of engaging across the UK and galvanising the research ecosystem.

I cannot say. This is partly experience and judgment. I would argue that the two things that are needed for ARIA to change the future of the UK and drive that success, based on what I am seeing, are a big enough portfolio to make sure that we have a chance to get lucky and have one hit, and enough time for them to mature. It is still early days, but that is the case we believe in.

Lord Drayson: Both of you have the commercial and venture capital experience to know that, unless there is the constraint to ensure that the early funders, the early founders, are not diluted out of sight in subsequent rounds, and given that for these projects to



HOUSE OF LORDS

break through they will need significant funding, how do you ensure—given the challenge the UK has, outside the EU now and not aligned with the US, for the moment at least—that that growth will happen here when that scaling investment has to come in?

Matt Clifford: ARIA has been designed very explicitly as a value-creation mechanism, rather than a value-capture one. That is the right call. What do we know about the benefits of innovation and publicly funded R&D? Innovators capture only about 2% of the value created. The vast majority of value created by R&D is captured by consumers. What really matters for economic growth is not equity value but adoption. It is absolutely important. In my other job as a venture capitalist, I am all about value capture. But for the UK to benefit what matters is not ARIA having a stake in the outcome but broad adoption of important technologies.

If it turns out that an ARIA programme ends up changing the way global pharma interfaces with the brain, yes, it would be great if there was £1 trillion company that captured that value in the UK, but the primary benefit to the UK would be its people having life-changing access to new drugs. That has to be the thesis of all publicly funded R&D. Value capture almost never justifies the cost of publicly funded R&D.

Lord Drayson: Will it drive economic growth?

Matt Clifford: It absolutely will. The main way in which technology drives economic growth is through adoption, not value capture.

Ilan Gur: I will share one last personal titbit on this front, because I can tell that this is important to you.

Lord Drayson: I am not convinced; 36 months of investment flowing out of the United Kingdom as a result of this problem.

Ilan Gur: ARIA is a start-up. Start-ups are a bet on people and timing. The bet here, as I mentioned earlier, is that through ARIA we can activate the UK research ecosystem in new ways. The timing is that right now is the moment to do that. I have said publicly in other fora that I would not have accepted this appointment five or 10 years ago, because I visited the UK and engaged with the brilliant scientists and engineers here, I did not see the same spark of wanting to be supported, activated and empowered in a different way.

The question I got from many others, including my former colleagues at DARPA or ARPA-E, was: will you be able to create an agency that has the right early culture and engages the right type of people with the right mindset? From my perspective, we have



HOUSE OF LORDS

nucleated something that feels right. At this moment, there is a chance for that to grow. That may not be satisfactory as a quantitative or data-based answer, but I do think there is a moment here.

Q25 Lord Borwick: This is a question for Matt Clifford. When we passed the ARIA legislation, a lot was made of the independence of its board. Has that proved to have been overplayed? Do you find any advantages or disadvantages, and is it real? Does your position in having a separate relationship with the Secretary of State whereby you are consultant to him confuse your independence as chairman of ARIA?

Matt Clifford: Thank you for the great questions. One thing has been extremely pleasing so far and beyond both our expectations is that when we were both appointed we worried a lot about independence and whether ARIA would really be given the freedoms laid out in the Act. It is pleasing, and it feels to us, that in DSIT not just the letter but the spirit of the ARIA Act has been embraced by the sponsorship team, the officials who we work with week to week, and the political leaders both before and after the election.

The commitment to fully delegate governance to the board and not attempt to run ARIA from 100 Parliament Street has been treated with the utmost respect, and we have had no issues whatever with that. As you rightly point out, I have worked in an unpaid and part-time capacity for DSIT on a number of occasions over the past few years. As the Prime Minister's representative for the AI safety summit last year and currently as the chair of the AI opportunities plan, I would like to think that this has been an asset for ARIA in deepening my relationships within the department, both at the official level and with the political leadership. That hopefully gives me credibility and strength of personal relationship to make the ARIA case. As Lord Drayson said, this will be very important in the months ahead as we push for ARIA's settlement at the next spending review.

Q26 Baroness Young of Old Scone: There was one point you made that struck me, which is that you are now out to recruit programme directors afresh. That is quite quick. Do these folk not hang around, and is there a risk to programmes by them sloping off?

Ilan Gur: I am so glad you asked. One of the core design features of ARIA, and this is one of the things that DARPA has inspired, is that our programme directors are recruited on a term basis. Each of them will bring fresh perspectives, intensively develop the thesis for a programme, and manage and direct that programme. Then



HOUSE OF LORDS

they will go back to other parts of the ecosystem. That exists for two reasons. One is that it allows for a purity of incentives. We never have to worry that any of our programme directors are making decisions in guiding these programmes based on their long-term career prospects within ARIA or what their boss wants. They are just focused on how they deliver.

Secondly, it ensures that we constantly are refreshing. Frontier technologies change. The assumptions that you made three years ago may no longer be valid. Your internal biases get fixed. It was always part of the plan that we would have a refreshing rotation of programme directors. We recognised two things, and this speaks to the reason we are recruiting right now. One is that even though we have seven incredible programmes, that is not enough of a portfolio to deliver on the mandate you have set. We know we need some new programme areas, and we want to ensure that we build in overlap between these programme director groups to drive learning and build that base.

Baroness Young of Old Scone: You have no doubt thought about it, but does that create any risks?

Ilan Gur: I do not think there are any unknown risks.

Baroness Young of Old Scone: What are the known ones then?

Ilan Gur: There are trade-offs. Let us put it that way. If you have people rotating out and fresh people coming in, the question is whether there is institutional knowledge that gets lost in that transition. Are there dislocations or disruptions that happen when someone new comes in to take over an area or shape something new? My experience from DARPA and those other agencies has allowed us to build in, pre-emptively, some of the best ways to mitigate those factors. But, all in all, the view is that the benefits far outweigh those downsides.

Lord Lucas: It strikes me that the scalable neural interfaces could do with a much wider ethical consultation than just something internal to you. The opportunities to use them to produce a well-ordered society on the North Korean model are fairly clear. If they are used in mental illness, what about using them in the classroom? It strikes me as a dangerous bit of territory.

Ilan Gur: Obviously, the work in that programme is very early, so you might think about those potential consequences down the road if our flare ends up driving further development. There are a few ways in which that is being built in. I mentioned that Cambridge University Health Partners is running alongside us to think about



HOUSE OF LORDS

some of these issues. We are recruiting a set of ethical and social responsibility advisers specifically to look at the work being done across projects in that programme. That programme is funding, alongside some of the technology innovations, work focused specifically on patient engagement and research, because we have realised that there is a dearth of knowledge of that. Those will be the initial approaches we take.

The Chair: Thank you very much for a fascinating session. It has been great to hear from you about the rich research and entrepreneurial base you find in the UK, which we hear a lot about, but it is great to hear that from you, coming from the States. I think you may also have picked up a concern among some of us about how we turn this into value creation for the UK, because we do not want it to be another story about all that rich research and entrepreneurial base being exploited overseas. One of the challenges ahead is how we keep that in the UK and make sure it delivers for UK economic growth. Thank you very much to both our witnesses.

Ilan Gur: On behalf of the whole team, everyone feels a great responsibility and privilege to serve the mission.

The Chair: Thank you.