



HOUSE OF COMMONS

Defence Sub-Committee

Oral evidence: Developing AI capacity and expertise in UK Defence, HC 429

Tuesday 20 February 2024

Ordered by the House of Commons to be published on 20 February 2024.

[Watch the meeting](#)

Members present: Mrs Emma Lewell-Buck (Chair); Sarah Atherton; Martin Docherty-Hughes; Richard Drax; Mr Mark Francois; Jesse Norman; Sir Jeremy Quin; Gavin Robinson; Derek Twigg.

Questions 1-36

Witnesses

I: James Black, Assistant Director of the Defence and Security research group at RAND Europe; Dr Simona Soare, Senior Lecturer at Lancaster University; and Air Marshal (Retd) Edward Stringer.

Written evidence from witnesses:

- [RAND Europe](#)

Examination of witnesses

Witnesses: James Black, Dr Simona Soare and Air Marshal Stringer.

Q1 **Chair:** Good morning, everyone. Welcome to our first session this morning on developing AI capacity and expertise in the UK. We are very pleased to be joined today by James Black, who is an assistant director in the defence and security research group at RAND Europe; Dr Simona Soare, senior lecturer in strategy and technology at Lancaster University; and, virtually, Air Marshal (Retd) Edward Stringer.

Before we move into our questions, can we have some comments from our experts?

James Black: I am James and I am assistant director of the defence and security group at RAND Europe. RAND Europe is the European arm of the RAND Corporation, which is a non-profit research organisation and the world's largest policy research organisation. It has been around for 75 years. It was in the "Oppenheimer" movie—I do not think we made it into the "Barbie" movie, but perhaps if there is a sequel. I will say up front that I had surgery on my hand yesterday, so if I wince at any point, it is a reflection on that and not on the line of questioning, I hope.

I am very interested in supporting the Sub-Committee today. I come at this from the angle of a lot of research that RAND has been doing for the UK Ministry of Defence and other allied Governments around the world. In particular, I have worked quite a lot in relation to the US and Australian Governments, thinking about things such as AUKUS, as an example.

Dr Simona Soare: Good morning. I am Simona Soare. I am a senior lecturer in strategy and technology with Lancaster University. I have recently joined Lancaster University. As you might know, the university is investing quite a lot in putting a new initiative in place to support the cyber corridor and to support the development of the digital ecosystem around Manchester.

My background is in defence studies and I am a strategic studies expert. Particularly out of that, I focus on defence innovation and the adoption of emerging technologies such as AI in defence applications. I have done a lot of research and have worked quite significantly with armed forces and Governments, in Europe mostly and the US, on proposing a new way of thinking about defence and how you generate and sustain military power. That is software-defined defence, which rests so much, obviously, on AI and digital transformation of defence.

I want to say that it is an absolute privilege to give oral evidence today to the Sub-Committee. Thank you so much.

Chair: You are very welcome. Over to you, Edward.

Air Marshal Stringer: Good morning. Thank you for allowing me to participate this morning. I retired two years ago. My last job was director



HOUSE OF COMMONS

general of joint force development. In 2018, we kicked off a programme called Information Advantage. After that, we produce other documents that are referenced in the defence AI strategy, which I am sure all the members of the panel have read, such as the integrated operating concept. That was actually written in 2019—it says “2020” here.

I am not going to say too much in opening, but what I would like to try to get across today is the sense of some form of calibration as to how fast we are moving. By analogy, to what extent is defence behaving like Amazon—that is, realising the real advantages of the information age and setting up a new business model—or to what extent is it doing a Waterstones, a traditional bookshop with a slightly better website attached?

I will finish by saying that, while I would not disagree with anything that is written in the defence AI strategy, I can see whole paragraphs that are lifted from what we wrote in 2018. That is six years ago now, and look at how quickly people have moved, not least the Ukrainians, with *necessitas mater inventum*. It is after we started thinking about this that Mykhailo Fedorov was put in charge of digitisation across the whole Ukrainian Executive, and look how quickly they have moved. Those would be the thoughts I would open with: to what extent are we being ambitious enough and are we moving quickly enough?

Q2 **Martin Docherty-Hughes:** Good morning to everyone; it is nice to see you all. The question I have been asked to ask, and will maybe expand on a little, is this: what are the main strengths and weaknesses of the UK’s defence AI sector? Maybe we will start from the nearest—so, James and then Dr Soare.

James Black: I will start with the big picture, so I do not eat too many of the other two speakers’ sandwiches. I think the UK has a genuine opportunity but, echoing the comments that the Air Marshal already made, it is a question of whether we can build on the strengths we already have at the pace we need to.

In terms of what those strengths are, you can look across the UK AI sector as a whole, or indeed UK tech as a whole, and certainly see some real strengths in things like basic research. We have a very strong academic and university sector, and always have. That is not just true when it comes to AI; there are also a number of related fields that feed into it. We have the City of London—in international terms, a comparatively successful financial hub—and lots of inward investment, venture capitalists, and so on.

We have also seen that, at least within Europe, the UK has retained an ability to attract top tech talent more than other European countries. We have to recognise that that is still a long way behind what Silicon Valley and the US can do in terms of attracting talent, whether from the UK itself, out of the UK or from other countries, where people are choosing to take up jobs in the United States rather than in London, Belfast or Manchester. We have some fundamental strengths in those sorts of areas.



HOUSE OF COMMONS

If you dive into areas where the AI sector has been very successful, we have done very well in things like natural language processing and computer vision. But we have also done well in adjacent areas that bring in some of the UK's strengths in things like the social sciences, the humanities and law—so, thinking a lot about AI ethics, AI governance and AI safety, about how you apply some of those principles on a technical level and a procedural level, and about spillovers between AI and other nearby sectors, such as fintech, biotech and so on.

All of those position the UK with a lot of quite good generic strengths. Of course, what you then need are the appropriate Government strategy, resource, implementation, investment and so on behind that, to then capitalise on them. I think we have to recognise that the UK's strengths have a half-life. These things are not static, so you have to continue to invest in them to retain them. We cannot take it for granted that we will have those strengths in five years' time. You also have to recognise that they are comparative, so other countries are investing as well to try and catch up in areas where we are particularly strong. I will hand over to Simona to elaborate.

Dr Simona Soare: Just to follow up on what James was just saying, it is important to highlight that one of the strengths in the system is the fact that it has also been significantly supported by funding—more so than our European partners and obviously on a different scale from our American colleagues. The UK ranks very well when it comes to funding of artificial intelligence, and particularly defence artificial intelligence, nearing now approximately £1 billion. That is an estimate that my team and I have put together. By comparison, that is twice as much as our French or German counterparts are investing in defence artificial intelligence. There is significant funding and that is, I think, a strength.

I also want to emphasise that there are important structural weaknesses in the way that the AI ecosystem is built in the UK. One of the things to point out is that there is no particular separate defence AI ecosystem in my mind. It is a general-purpose technology, which means that it is developed mostly by the private sector, and so the chances of attracting AI applications in defence rely a lot on the health of the AI ecosystem writ large. Here, we have some interesting dynamics.

First, it is still a maturing ecosystem. We have under 3,500 companies and the vast majority of those, to the tune of 75% to 80%, are small enterprises and start-ups. The challenge with the structure of this ecosystem is that, obviously, the turnover of start-ups in particular is very high: about one in five or six make it past the four-year timeline of their lives. That is very important, especially when we are thinking about defence applications.

The second aspect as to how the ecosystem is built is that it is very concentrated. Eighty per cent of the companies are located in the south and south-east and nothing much outside. With the cyber corridor, there is hope that the digital ecosystem will develop in the north, around Newcastle where we have the data centre, but that is work that needs to



HOUSE OF COMMONS

be supported by the Government and the MoD so that you can actually bring the providers closer to elements of the force.

Last but not least, I would like to emphasise the fact that with the structure of the current ecosystem in the UK, it is important to highlight that, should we need to increase capacity for artificial intelligence applications in defence, there would be a very thin layer of resilience in the ecosystem. The reason for that is, again, that high turnover and the uncertainty that the small and start-up companies face on the market.

There are important elements of strength, which have been highlighted. There are also important vulnerabilities that we have to take into consideration.

Martin Docherty-Hughes: Thank you. Edward, do you have anything to add to that?

Air Marshal Stringer: I agree with everything that has been said. By analogy, I would point to the car industry and say that we are Formula One but we do not have Toyota or Tesla, which I think has been said, when we are looking at scale. Also, talking around friends in big industries who are looking to really invest and have invested, and are trying to tie in with universities, it is difficult to find UK students in depth. In fact, you will find a lot of Chinese students, but you will not find so many Brits. There is more to do there and there are experts on this panel from the university sector who might add more colour to that than I can. Across defence AI, you have some very good people working very hard, but within a slightly fragmented system. I agree with that comment about the AI ecosystem.

My final point, and this is expressly about defence, is that, under Levene, what you actually have is each of the services being allowed to go off and do its own thing. On the advantages of scale, when one thinks about digitisation, we always have to start with thinking about the data and therefore shared data and lots of data, but if you start developing your own little programmes and stovepipes, you go away from that, so you are putting some digital baubles on an analogue Christmas tree. I still just get a sense that defence hasn't gripped the real possibilities of this and digitised across the whole of the armed forces. But I could point you to lots of really good people working very hard in small but, I'd argue, slightly over-siloed areas.

Q3 **Martin Docherty-Hughes:** Thank you. I do like the analogy of an analogue Christmas tree—we might take that further at some point.

To go into a little more depth on some of the issues that you have raised, do the UK Government have the data they need to properly understand the nature of the UK's defence AI sector? Would you like to come in here, James? I know that you have mentioned this.

James Black: I am going to very quickly jump in and say that that challenge isn't unique to the AI sector. The UK Ministry of Defence in particular has been taking a number of initiatives over the last few years to try to improve its understanding of more traditional defence suppliers



and SMEs. Going back a few decades, the Ministry of Defence has tried to shift a lot of the responsibility and, indeed, the risk for managing supply chains on to prime contractors contractually, rather than having to do that itself. That has made the lower levels of the supply chain more opaque to Government in general. That's true across a variety of sectors. The MoD has been trying to find new ways to collect data and, crucially, to incentivise companies to provide data. I think that's another point that we often miss when we start talking about collecting data on the AI sector. As has been mentioned, given the size of some of these start-ups, the turnover, the uncertainty they face in their financial futures, they need to have the incentives to provide the data that you need to track how they are doing.

Q4 **Martin Docherty-Hughes:** Do you think that part of the problem is that we as legislators and the policy makers in the Department who assist us aren't really that up to date on what AI actually is? I declare a non-pecuniary interest as chair of the all-party parliamentary group on Estonia. If I were to talk to my colleagues in the Estonian Parliament, they would certainly be able to give me a more profound sense of what AI and cyber capability is required not only by their Defence Department but by the entire state.

James Black: To come back very quickly on that—sorry, I can see the time—I certainly agree that there's a cultural issue there. If you look at some of the smaller countries that are disproportionately good at AI, certainly in relation to the military, places like Israel, Estonia and Finland are countries that have clear existential threats and a long-standing culture of total defence, comprehensive security, whole of Government, whole of society—whatever kind of language they use. But it is something that is much more embedded. You have people who are reservists—pretty much every adult male and, increasingly, many adult females—so it's a very different culture around understanding risk, security threats and defence in general. Obviously, countries have also recognised that, because of their size, they can't really use mass to solve their problems, so they have to use technological sophistication, which is why they have been very good at investing in digital technologies such as AI.

Dr Simona Soare: To go back to the first question, I believe that there is data, both in the MoD and in the governmental system, the whole-of-Government enterprise. However, there is a very different conversation about where that data sits, what is being done with it, who knows about it and who is exploiting it. Over the past five or six years, the state apparatus has come a significant way in trying to understand what data it needs to drive strategy implementation. On the digital side, though, I feel we are still scratching the surface.

I will give you an example. If I look at the defence and security industrial strategy, which should give us a good idea as to how the defence ecosystem is thinking about its relationship and supporting the defence industrial ecosystem in the UK, I find very little there in the way of digital, in the way of AI. There are elements that are clearly signalled for prioritisation, but there is not much in the sense of data related to the



current size and structure of the ecosystem, which tells me that perhaps that might be a little bit of a shortcoming in terms of data.

Another example is to do with the innovation agencies. It is very clear that the MoD is trying to reach out as much as possible to the private sector, to understand who are the newcomers on the market, and who are the new start-ups that do exciting things in technology and AI. However, when we have done research on this topic—we have done research in the UK as well as across Europe—we have asked governmental agencies to tell us what is the process of collecting that data. More usually than not, the data sits in an Excel database.

When we find out about a new start-up that does exciting AI technologies, we put it in our little Excel spreadsheet. It sits there until we have an opportunity to engage with the company around the specific type of program or Word system, and so on. That can take months. These are start-ups; they need at least £25,000 a month to keep afloat. If it takes the MoD five months to reach out to them, they are out of business by then, quite likely.

Once it comes in, the use of the data is also problematic. It is not just a shortcoming in the level and quality of the data. It is also how to use it once you get your hands on it. Related to the legislators and decision makers, and the level of AI readiness there, that is an accumulation. It is a cumulative process, and nobody is expected to be a software engineer overnight, or to fully understand it. It is important, however, to start to dig at how AI can be embedded into military applications. Parliament has definitely taken upon itself to make strides towards that.

We also have a luxury that is perhaps not available to some of our partners and allies, such as Estonia, as James mentioned, or Ukraine. Ukraine is on an accelerated learning curve, including for its policy makers, because it needed to be. The UK does not have quite the same pressures.

Martin Docherty-Hughes: Briefly, Edward, do you want to come in on that point?

Air Marshal Stringer: I would just say that a lot of companies find it very difficult dealing with Government and with the MoD. The reason is speed. I quote an ex-military friend from last year. He set up a vibrant company, doing very well, expanding into the US. He said, "I thought it might take me a couple of years to beat down the door and get the MoD to understand. It has taken over four." That is the problem, and supports what you just heard about start-ups having to move quickly. Everything is moving very fast in the world of AI, and the machinery—the wheels of government—are not keeping up.

Q5 **Jesse Norman:** The questions I have are of rather an operational character, so in the first instance I will address them to Ed Stringer, and if others in the room want to come in, that will be great. There is a paradox here. We have been using machine learning in defence for a while, but no one had heard of a large language model two years ago.



HOUSE OF COMMONS

There is massive change going on at the moment. If you are in that sector in the civilian side, it is like a gold rush. You are not going to university because you can do better, faster in this game-changing world, just by, as it were, directly investing.

That picks up on the point Ed Stringer was making. So, Ed, can I just ask you a little bit about the differences with current defence, as it is managed and procured, and could you expand a bit more about the model problem that you described? I suspect we are just way off the pace on this. With all the good work that is being done, we probably need to recognise that.

Air Marshal Stringer: Yes, we need to get back to the psychology of contracting. We are still configured to go to big primes and beat out a contract, and put in thousands of pages all the cardinal points. Then our lawyers and their lawyers will argue over the 20 years of the programme, before a big bit of plant is delivered. Software companies do not work that way. They use the minimum viable product method. They want to work with you on a problem, get something that works, get it out there, test it, get the data back and keep evolving it. It is very difficult with our own psychology of contracting to write an old traditional contract for that. What you will hear from companies is something about, "You work with the MoD and they will demand yesterday's technology tomorrow, and then they will delay it and then they will haggle over the price."

Certainly some companies have ended up pulling out of the market or have hived off their defence element into a sort of bad bank within the company because they just cannot move quick enough and they do not want to drag everyone else down. That is not unique to defence. It is a problem across Government. I believe there was a unit in the Treasury set up around the time the integrated review kicked off to have a look at some of this. I do not know what has come of that. I think it might have drifted away a little bit, but if we could do something to both help Government and help industry, it would be thinking through exactly how you contract for software-driven products.

I will just finish with a simple analogy. I am sure most of us around this panel will use a Microsoft or Apple product. Do you know when Apple is going to drop a patch to update your systems? Do you know if it is going to come this week, next week? You might get two in a month; you might not get one for three months. That would not work with an MoD system where you want to know—I want a drop in May; I want another one in July; I want to know what you are going to give me in September. Well, they can't know. Everything is evolving and moving quickly, but you've both got a vested interest in the computer working. You both want it to be secure.

I will walk away from Apple and go somewhere else if they do not keep on top of things in cutting edge. If they do not keep on top of things in cutting edge, they are not going to make money. If there is one thing that could come out of this, it would be to think through how you work with tech start-ups, especially in AI, where things are moving so quickly that



you end up with genuine benefits for both. At the moment, we are not getting benefits for either.

Q6 **Jesse Norman:** That is very helpful. Chair, we might want to think about picking up the point about the Treasury team that has been raised by Ed and maybe pressing that question. I don't know, James or Simona, whether you want to come in on the back of that. Do you concur with the analysis? Have you seen any evidence that this kind of thinking is now being implemented? Obviously, there is a gigantic bureaucracy all sitting there waiting and hoping to keep their jobs and writing documents in the usual way, which would be threatened by a more rapid form of procurement.

James Black: Absolutely. I have just finished a couple of studies, one for a US congressional commission on defence planning and budgeting reform, which is about as exciting as it sounds, but is very pertinent to this. The second is specifically on AUKUS, more rapid innovation and the development and fielding of AI-related capability. I think that what we are seeing, as Ed has said, is that this challenges the model across the entire capability life cycle. It challenges what it is that we are trying to procure in the first place. You are no longer buying a big lump of metal—a ship, a tank, an aircraft or whatever—and doing it in a 10, 15-year timeframe and having most of the value wrapped up in hardware. You are instead now looking at a much more rapidly evolving force of unmanned systems, as well as crewed systems, much more emphasis on connectivity data, AI, et cetera, and then consequently on software. What you are trying to buy is different.

As has been mentioned already in relation to Levene, that forces you towards much more of an integrated approach across the services, across domains and across the top-level budget holders. That is really tough on a bureaucratic and cultural level to achieve, because clearly every different two star, three star, four star has their pet thing that they want to invest in and the requirements they are focused on. It then changes how you are buying it, as has been mentioned.

You need a remarkably different approach to how you develop software-related capability. It is not the waterfall model that you need. There are lots of good examples of things like agile and DevSecOps and other things being implemented in small, siloed areas across defence, and that is good. But we are not doing that at scale, when you have had decades of more traditional procurement approaches, slower procurement approaches, and, crucially, risk aversion. Because obviously defence has a remarkably different attitude to risk than somewhere like Google, you cannot move fast and break things in defence in quite the same way, or you do not want to because of all of the safety concerns and political concerns.

Finally, it affects who you are buying from and the influence you have over them. You are no longer just talking to the BAE Systems and Lockheed Martins of the world. You are talking to potentially much smaller, innovative start-ups focussed on AI. Although some of those, particularly in the US, emerged with a focus on the defence sector as their primary



HOUSE OF COMMONS

purpose and customer, for many defence is a long way down their list of customer priorities, so the level of influence that the MoD has as a spender is more limited. The MoD can influence the shape of somewhere like BAE Systems because it is almost a monopoly/monopsony situation in some areas, so it has a lot of leverage financially. You can't do that in AI when the company is selling to a general audience and, as you say, doing it much quicker; in contrast, defence is behind the curve.

Q7 Jesse Norman: Thanks. I am keen to come to Simona in a second, but can I ask a follow-up question? We are not a million miles away from the point you made becoming a reality, in the sense that all the value is in the software and the compute that supports the software, and therefore it is really only being controlled by a handful of companies—four or five around the world. The question for a defence project is, “With which one are you allying, in the knowledge that by the time your programme or platform is ready it may be someone else in the driving seat?” How do you solve that problem?

James Black: There are some practical issues about how you contract and field a capability, and then there are some broader industrial policy issues. Let me start with the latter. On the industrial policy side, if we go back to the Defence and Security Industrial Strategy, there is a real question about the “own-collaborate-access” model—the idea that the UK should have certain capabilities where it needs sovereign control, others where it can work with allies and partners and others where it can accept that it is just going to rely on the open market, with the risk that comes with that. Defining that in relation to hardware is tough but measurable. It gets harder in relation to the software, for all the reasons we have already talked about, particularly given how quickly things move.

You have a different relationship with your suppliers anyway. It is not that they are providing you with a vehicle, a missile or whatever, and then they hand over the keys and, although they provide rolling maintenance, you largely own and operate that capability, so it is more of a transactional relationship. When you start talking about software delivery over time, particularly as a service, you are talking about a much more intimate relationship between the user and the designers. You are providing constant feedback, data and so on, so it is much more enmeshed.

Contractually, there are things you can do. Promoting open systems architecture avoids vendor lock-in, so there are good practices there, but you still have to find a way of incentivising companies to work with you and collaborate in the direction you want, including around things like IP and those sorts of issues.

Q8 Jesse Norman: That's brilliant. Picking up on that, if you want to do object recognition—for example, target assessment—can I put it to you that you would be mad to go to anyone other than a 27-year-old working in a small company specialising in this area, which has probably been thinking about mammograms but might be thinking about weapons targets? You wouldn't necessarily go to one of the big primes, because they don't have the expertise and don't have the culture that will attract



HOUSE OF COMMONS

those people. Is that right? If you are a big prime, how do you manage the demands of the sector, in terms of the challenges offered by AI?

Dr Simona Soare: The big prime sector is built for a very different business model, which is far closer to what the procurement system is in the MoD—they feed off each other. You are right that when you want to get cutting-edge technology today, you are more likely to find it with a 27-year-old than with a prime, although that is not to say that primes are not technologically advanced. This is something that they recognise themselves.

In my opinion, if we look at dynamics in the market, particularly around artificial intelligence, we see a number of very interesting partnerships between the primes themselves and newcomers on the market, which can be start-ups but are usually small and medium-sized companies. They have a very different business model. They don't need 200 software engineers in one company to work on one product; they will hire three, who will be the best on the market. The company will make that financial and capital investment, in relation to a big prime or an accelerator, in order to drive them forward and put them on the radar for defence. Those are two very different systems and business models.

The MoD is still not fully leveraging those market dynamics. It is very difficult, even for a bureaucracy as big as the MoD, to engage individually with start-ups. That takes a lot of time and a lot of human capital. If we at least start by capitalising more and leveraging these new market dynamics, there will be a gain in attracting new technologies faster.

There will also be a gain because, whatever we are deploying, it will still have to function as part of a bigger system and platform. It will not just be in isolation so you cannot treat the existing inventory as a complete vacuum. It does not work that way. Therefore, it still needs to work on an industry-to-industry basis as well. I think that that is one of the areas where we can improve the situation faster: by leveraging market dynamics fully.

Jesse Norman: So you would be pushing us more towards something like a life sciences relationship with the pharma sector, where the really smart stuff is getting developed at the cutting edge in lots of the smaller companies, but there are established patterns of bringing that technology into play.

Dr Simona Soare: Absolutely. It is about sending a demand signal as to what the MoD needs in terms of technology and what problems it needs to use that technology for, and then incentivising market dynamics in that direction, rather than being the arbiter.

Jesse Norman: Thank you. Ed wants to come in on this as well, and then that will be it from me.

Air Marshal Stringer: I was going to pick up on that relationship with the primes and also give some evidence. In 2020, to get around that problem of the analogue Christmas tree with a few digital baubles, we tried to write



HOUSE OF COMMONS

a blueprint that looked across all five domains, and therefore all three services and all four top-level budgets. We looked forward and said, “This will be demonstrated in 2024 and Exercise Steadfast Defender.” Isn’t it amazing how quickly four years goes?

Having written it, we sent it around to quite a few of the big primes, both defence and IT—those getting into cloud and so on—and also to a heck of a lot of the smaller digital companies. I had better not name any of them, because I will be putting words into their mouths in the public domain. There was an absolute consistency in what came back: that this was doable, that the technology was pretty much there now, and that, given four years, there was certainly very little risk. The big point here is that the SMEs liked it, because they could see where they could dive in and add advantage.

We ended up having four hours with the CTO and CIO of one of the big defence primes. They said, “This is great, because it allows us to do what we really want to do, which is not chase stovepiped individual platforms. This is a call for how you integrate across all five domains, and that is what we are very good at. We can help the SMEs to scale.”

The final point I would add is this. Everyone talks about working with SMEs, but we call them SMEs because they are small. It would be difficult for them to scale rapidly and deliver a contract across the whole of defence or even across the whole of Government, but the primes can help them do that. So, yes, I think pharma is a good analogy here. When we talked about this in defence, one of the big consultancies that I advise immediately saw parallels in pharma. I think it is valid.

Chair: Thank you very much. Gavin Robinson is going to explore the Government’s defence AI policy.

Q9 **Gavin Robinson:** Good morning to you all. Mr Stringer, I am going to stay with you. I think that a number of your analogies have been illuminating for us, and I think they have chimed with members of the Committee. You were director general of joint force operations in the Ministry of Defence, and you have highlighted a number of issues that you were involved in particularly and that are there just now. Did you leave the Ministry of Defence with a sense that it had grasped and was rising to the challenge of AI and future development, or did you leave the Ministry of Defence rather frustrated?

Air Marshal Stringer: I would have to say that I left frustrated. I left frustrated by what I sensed was, in a phrase nicked from a good American general I know, a “say-do gap”. I think we got to the point where lots of people could trot out buzzwords around defence in AI, but they were trying to put those as bumper stickers on the cherished programs they had been running for a while and wanted to push through. Structurally, the post-Levene system incentivised the services to do that.

I think this is getting better now, but if you go back just a couple of years, one of the bits of work we did was using quite simple machine-learning



HOUSE OF COMMONS

tools to work out who was doing what in experimentation. We found out that, in just one service headquarters, two one-stars, who were personal friends, were working on exactly the same problem and did not know it. Across all four TLBs, you had people racing to do almost the same thing with data and often employing the same companies, but in quite a stovepiped way: looking to solve small-scale problems because that is what they are incentivised to do.

Of course, the service chiefs are incentivised to look after their own institutions, which becomes ever so slightly competitive. I am sure we will see that next time there is a spending round. I think that structurally there are problems here, and they are especially acute when you deal in the fast-moving world of high tech, where it only works if you share data across a whole enterprise.

- Q10 **Gavin Robinson:** Thank you very much. I am glad I asked that additional question. I will ask the next question of you, Mr Stringer, and then widen it out to your colleagues. The defence AI strategy published 18 months ago—I think this attaches to some of the points you made—talked about future AI. However, we are just coming up to the second anniversary of Russia’s incursion into Ukraine, and aspects that were described as future AI are now being developed and deployed on quite a scale. From your perspective, does that defence AI strategy urgently need a refresh?

Air Marshal Stringer: If you were to refresh it, you would always be behind the drag curve. Things are changing month on month in this world. To put that into focus, I was talking to a good friend who is absolutely an international-level guru on AI at a conference in May last year. He could say then that training a model would cost you \$10 million in January this year; now, already, that has come down to \$1,000. I rang him up because I was speaking at a conference two weeks later and wanted to check my facts. I said, “I must have got these figures wrong, Robert. Where is it now?” He said, “I’ve just opened a paper from Stanford. It’s down to \$600”—and it could have been done a lot more quickly.

When things are moving at that sort of speed, you cannot keep up if you write strategies. I was rereading the strategy only yesterday, ahead of this session, and while you cannot disagree with any of the aspirations in there, I would point out that the data strategy it references positively is from 2021. We started a thing called Information Advantage in April 2018; at the very first meeting, we identified data as the crucial commodity and that we needed to very rapidly come up with a plan for how we were going to store it and access it—we were pointing at a military cloud or combat cloud for doing that. Well, it is three years to even write a strategy for what we are going to do with data. There are classified elements here, but I know that the intelligence agencies were behind us in our thinking back in 2018, 2019 and have now moved much quicker. So it can be done across Government.

I would say that it is not a question of refreshing the strategy; it is a question of revisiting how you do procurement and product development in the information age, because if you stick with writing strategies,



publishing them and trying to influence things, you are always going to be behind the drag curve in this world.

- Q11 **Gavin Robinson:** Thank you. That sets things up very nicely for our other two witnesses. There is a recurring theme of criticism around strategies getting to the point of publication and then being out of date quite quickly, or there being this gap between rhetoric and delivery. Do you need to have an infrastructure or an architecture that allows you to evolve within it, rather than a set piece that is solidified and foundational and that fails to move with the times and the pace of change?

James Black: I think that that was really the intent with the Defence AI Strategy—to provide something that would be a bit more long-lasting than your average strategy document out of that Department or other Departments. That was reflected in a desire to, first of all, address up front some of the ethical considerations, which are obviously most acute around defence, but which also apply in other sectors. It was recognised that that is about building trust and buy-in to the strategy from audiences within defence, across Government, the general public and politicians, and, of course, internationally. So it was about the sorts of things that were designed to last.

Then, the strategy focused a lot on some structural, cultural and process changes within the Department. You had the Defence AI and Autonomy Unit, but also the setting up of the Defence AI Centre. Then, within that, there were initiatives around things like CommercialX and trying to go after some of the more novel contracting approaches and so on that the Air Marshal has just alluded to. Similarly, it was thinking about things like career management, career structures, the skills for AI that are needed within the Ministry of Defence from civil servants, military personnel and reservists, accessing industry, and putting in place a lot of other basic enablers. I think all of those are very positive in terms of a move in the right direction.

The question now is the follow-through. It is both the scale of the investment and resourcing of this and the willingness to implement the sweeping level of change, which all three of us have alluded to, to really deliver on what has been said—to bridge that “say-do gap”.

As has been mentioned, a lot of very good initiatives have sprung out of the strategy. The question now is: can you follow through on them? Can you cohere and scale them? Because you can go back to any strategy that the Ministry of Defence, or pretty much any Government Department, has had over the last however many years, and find this life cycle of strategy, initial flurry of implementation, set up a bunch of new teams, appoint a bunch of new senior people, have a bunch of photo ops, and then the next strategy or departmental initiative comes around and all those fall by the wayside and are replaced by something else. We need to really follow through with the nitty-gritty of change management, and that can easily be distracted by a revised strategy and policymaking cycle. It is more about evolving what we’ve got and following through on implementation. There is not a lack of strategy—it is now, “Can we actually deliver it?”



Dr Simona Soare: There are elements of the strategy that can be worked with and continue to be worked with. It is a strategy—it will need a refresh at some point in time. That is how strategies work and how Governments work—there is a cycle. One area I would focus on is better understanding where the MoD is in relation to the implementation plan. The strategy comes with an implementation plan. It focuses on giving central guidance to the force and agencies, but also allowing them to decentralise implementation and build up their own AI capability plans.

However, I do not know exactly. There is not enough transparency on where that implementation stands today. It would be much more useful to start the conversation there, and then see whether there are elements of the strategy that need to be upgraded. Our colleagues across the Atlantic have done so. But again, the measure was: where are we with implementation, and what do we need to adapt or adjust based on the progress we have made with implementation so far?

There are elements of the strategy that are missing and would need to be addressed in a future iteration. One is clear metrics for success. The current strategy has none. One would assume that is because the implementation plan would follow, but that is not public—so again, we are going back to lack of transparency. The second is the fact that there is no exact timeline for us to better understand important adjunct issues such as progress on skills that are necessary for defence in order to progress the AI and digital transformation agendas. Again, there is not a lot happening there.

One of the priorities in the strategy was to support the development of enablers. As Air Marshal Stringer alluded to, that part is also falling behind. There are elements of the strategy that need to be reinforced in a future iteration—but does it need to happen now? I think implementation should take priority.

Q12 **Gavin Robinson:** I take it that there is a question around the metrics used, or how we are measuring or scoring. If an implementation plan is saying you must do x by a certain date, and you can greenlight that and say you have achieved x by a certain date, that takes no account of the need for evolution within the overarching strategy or how the implementation plan can be adjusted to account for evolution, challenge and the need to pivot, should that be where development and technology is going. Is that fair synopsis?

Dr Simona Soare: I cannot speak to the particular aspects of the implementation plan because, again, I am not privy to that document. But, reading the text of the strategy itself, it is very clear that the only progress metrics that we are given in the text are bureaucratic in and of themselves: “One year on, we are going to have the implementation plan”; “One year on, we are going to have the AI readiness guidelines”; and so on and so forth. These do not necessarily measure implementation towards achieving operational capability in artificial intelligence. They are just bureaucratic measures for where we are in the process.



James Black: I am involved in a lot of RAND's work around supporting strategy, making strategy and implementation in Government. When it comes to monitoring and evaluating strategy, we often talk about a logic model. You can look at inputs, activities, outputs, or outcomes, and you can try to measure each of those four things. Where traditionally the Ministry of Defence in general, not just in relation to AI, has often been most comfortable is in trying to measure inputs, activities and maybe outputs. We are good at measuring how much money has been spent, how many people are in a team, the fact that the team has been set up and so on. You can measure that it is doing some "stuff".

We are less good at making up that gap from outputs to outcomes. Is it delivering the right "stuff"? Is it having an effect, and is that robust in the face of changes in the external environment? We might have come up with a great plan and implemented it and be pumping out the outputs that we want to pump out, but if things have moved on in the wider world in terms of the threat we face and the technology available, that may no longer have the desired effect. To the extent that we can, we really want to try to focus on being relentlessly driven by outcomes, rather than those other things. Otherwise, as the Air Marshal alluded to, you end up in a situation in which people are tacking the language of AI on to their programme so that when they go into their meeting, they can justify the spending on their programme by saying "Well, it's desperately important for AI," and it is ticking these narrowly focused metrics that you have, but it is not actually achieving the broader strategic effect or the cultural change in the Department that you want to have.

Quickly, on that point, allied to that you need resilience metrics. That is something that defence has looked at in other areas of contracting. You reward your supplier in terms of whether they achieve the outcomes, but you also build in a few resilience metrics basically to ensure that they are not gaming the system. You do not want someone to really optimise the way in which they do something to hit your metrics and tick the boxes that you need them to tick when actually it is just building up loads of risk for the future. You get really good scores on your dashboard in year one, but that is because you have pushed a whole load of problems into year two, and it will be a nasty surprise for whoever succeeds you in post.

You can do things to ensure that you are achieving your outcomes but also check that you are not storing up risk for the future or for whatever poor soul takes your job on after you. Getting that right is tough, and doing that obviously goes beyond the gift of the people in charge of the AI strategy. It is also about broader strategy implementation, programme management and governance and so on across defence, which is not just the AI team's responsibility.

- Q13 **Derek Twigg:** I was thinking while you were talking in the session this morning. If you look at Ukraine, boots on the ground and tanks and so forth are still very important, but in terms of AI—a new way of fighting war—it has brought home to a lot of people the changes that are taking place, which, as you say, are on almost a daily or weekly basis. I wanted



HOUSE OF COMMONS

to hear about this, given all your experience and looking from the outside. You talked about how you have a strategy. We heard from Edward that you just have to keep on top and have to be moving things, and it is almost a daily change, so having a strategy in itself is not good enough. Do you think that the leadership understands where we are at in terms of AI and where we could go? I am talking about leadership in terms of politicians and Ministers, but also the heads of the services. Do you think that at that senior level there is a clear grasp now of what is opening up in terms of AI, where they are going to be, where they are going with it and what needs to be done? Do you have any grasp of that?

Air Marshal Stringer: Thank you for posing that question with an ex-MinDP sat to your right. I refer back to the say-do gap. As I said, I was looking through the strategy to ensure that I was up to date, and I noticed that it talks about how we must educate and train our senior leadership and that everyone must understand how AI works. This was published in 2022. I know that when I briefed information advantage to the senior strategic body in the MoD, chaired by the PUS and the Chief of Defence Staff, back in early 2019, that was a recommendation. Indeed, the executive committee went off down to the Defence Academy to take a course that we had actually been running for a couple of years by that point. Something is being written here that is already five years after the fact.

As I say, things are moving so quickly in AI. Even those who are more expert—coders and software writers—will tell you, when you really grill them, that they are not quite sure how AI works at the moment anyway. They are not quite sure how some of the black boxes, when you get into neural networks and other bits and pieces, are really doing their stuff. So it is very difficult to keep on top of it.

I would say here, to link to the previous answer, which I would agree with, that it is so much easier simply to measure inputs. “We have written a strategy. We have done this”—left-to-right thinking; just keep churning the normal processes that you are used to, with a few AI words added. You really need to do some right-to-left thinking. The difference between us and Ukraine is that Ukraine has a clear and present danger to deal with, and therefore it absolutely does need an output. If you look across defence at the moment, who is actually responsible for pulling together the three services to fight across all five domains in a way that actually is going to win the next war? It is hard to define who that person is.

The chiefs are actually responsible for the inputs, and we have already discussed that, so I will not rehearse that argument. They are responsible for the inputs—polishing the service that they would like to be. We have already seen this manifest in the fact that, of course, weapons stockpiles only come into consideration when you hit day one of the war. There is a structural problem here, in that no one is particularly incentivised to do this linking tissue between all the three services—the four frontline commands—and keep on top of it, because it is not the sort of thing that you measure day to day, but it is the sort of thing that you will suddenly find yourself wanting, come day one of the actual war.



HOUSE OF COMMONS

James Black: I would agree with that. The National Audit Office did a study about 12 years ago on the implementation of agile across Government—software development approaches and the mentality of how Government was doing its business. They came up with a framework for measuring where you were along a spectrum of “terrible” to “very good”. There was a quote for where we sit now, which is probably about 30% of the way along that spectrum: “Success comes about because of cumulative heroic individual success by people or by teams, and in face of systemic barriers.” So it is not something where we succeed because of the system; we succeed in spite of the system.

That is probably the situation we are in right now. We have some very good senior leaders and more mid-level and junior people who get it. They understand, I think, the urgency of the problem. They don’t, as has been alluded to, necessarily have to be the greatest technical gurus themselves. They understand, at least, the implications of this technology for their role and for the mission, but they are doing it in siloes, and they do not necessarily have the broader supporting structures of incentives and culture, career management and promotion—all those sorts of things—to drive that across the whole enterprise.

What we obviously see in somewhere like Ukraine is a very clear and present danger, which, because of the sheer urgency, cuts through a lot of the risk aversion that is baked into so many of our approaches to policymaking, procurement and so many other things.

There is a real human predilection towards worrying more about risks of commission than omission. We worry far more about sticking our neck out, making a decision and it is the wrong decision—“I have spent £100 million and it didn’t work,” or whatever—and getting called up in front of your boss or the Minister, and being chewed out and fired, and that is the end of your career, right? We care far more about that than the risk of omission—of quietly not doing anything particularly innovative, getting on and allowing the UK’s advantage militarily to slowly slip away quietly into the night—but you won’t be held directly responsible for it.

In somewhere like Ukraine, they don’t care about the risks of commission because they need to try something. If they don’t try something, they definitely lose. That brings us to that final word of mine, which is regret. That is where we really need to understand: can we achieve the sorts of change in a pre-war setting, as the UK is now, that the Ukrainians are achieving in a wartime setting? Can we do that when the attitude to risk is different and our approach to regret is different? The Ukrainians know that if they do not try something, they lose their country. They are willing to try. We know that life maybe goes on for five, six, 12 years or whatever, and it is more about people’s individual careers.

There are good examples of where these sorts of things have been done, and done well. A quick example would be something like the United States Special Operations Command, USSOCOM. They have a very different attitude to procuring urgent capability requirements. They actually sit the warfighter alongside the procurement manager, and the procurement



HOUSE OF COMMONS

manager knows you need to buy this thing really quickly because Dave is going on a mission in two weeks' time and Dave is going to get shot at if you don't do it right. That drives a fundamentally different attitude to risk. It is focused much more on mission success, rather than on bureaucratic tick-box exercises and compliance. That can be done at a small scale, but scaling that is really tough, and that is where we need leadership.

Dr Simona Soare: And it is not happening. To follow on from that, I fully agree with what has been said. One element that I want to bring to the discussion is the fact that leadership is very important. It is irreplaceable when it comes to setting the exact level of ambition—the benchmarks that we need to hit—so that everything else, from a bureaucratic point of view, falls into line. That is the importance of strategy, and leadership is irreplaceable. You don't need to be a software engineer; you need to understand the market dynamics and what you are relying on, in terms of producing this capability. That is tied much more to the cultural, organisational adaptation than to anything else.

We have spoken about procurement and regret. It is much more important to think of wasted time in procurement that materialises as capability lost than of something that just didn't go right in terms of procurement. That is capability lost not just because we are not able to utilise and leverage the technology, but because others—potential adversaries—might get to it and, with very cheap demonstrator technology, understand what they can do. They have no interest in growing the company and in the UK having a mature and self-sufficient AI ecosystem, so that capability loss is significant on multiple levels. When I think of the quality of leadership in artificial intelligence, I think about that level of ambition and about putting in place the strategy benchmarks—where do I need to be and how do I measure my success?

I would say, if I may be so bold—I am sure the Air Marshal is much better placed to talk about this—that when you meet the frontline command, the first questions in terms of leadership are, “Where are you with the implementation of AI? How many real AI projects are you driving? What is the progress there? What do you need in terms of funding? How flexible is the funding for you? How close are you to actual implementable capability?”

- Q14 **Martin Docherty-Hughes:** On that point, Dr Soare, you mentioned leadership. I don't want to labour the point about Estonia again, but when it regained its independence from the Soviet Union, it became a digital state because digital leadership followed the money. It was a Finance Minister who did that. Every Finance Minister since independence in Estonia has also been the digital Minister, and they have led political change. Do you think that, given that the Treasury holds the purse strings, it should listen to people like you and Edward, get a grip and understand the importance of things like AI and the digitisation of the state, not just in the MoD but in every other element? In Estonia, the Finance Department and the Chancellery started that transformation and everybody else followed suit.



Dr Simona Soare: I look at it from the outside. All I can talk about is the data that I have access to. Some of it is data that the Government, including the Treasury, is putting forward, and some of it is estimates based on the number of programmes that the MoD is churning out in AI. The data shows that the UK is spending quite significantly on digital capabilities and artificial intelligence. It is spending significantly more than its European partners. It is spending less than the US, but obviously that is a very different ball game.

Q15 Martin Docherty-Hughes: You said earlier that it wasn't having a real-term impact. That is my point: there is a lack of political leadership at the top. All that money is flowing, but it is not hitting the capabilities that Edward, James and you have been talking about.

Dr Simona Soare: Understanding how that budget is spent is key, in my view. Looking at the details is very important. The UK spends £4.5 billion on digital capabilities; almost £1 billion of that is AI. That is not small change in spending, but a lot of that budget goes into maintenance and upgrades of existing systems. That is different from countries like Estonia, as you mentioned, where building up that capacity and those capabilities is much fresher, so a lot of the budget is going into generating new capabilities. That comes with already having a partially digitalised force. That needs to be maintained. Understanding what the money is being spent on and trying to find efficiencies within that—not just in the sense of increasing it, but in streamlining the budget structure—is important, and that is an element for which leadership is important as well.

James Black: Governments around the world are grappling with how the machinery of government needs to change, as well as how the finance of government needs to change in order to seize the scale of opportunity and of threat that AI presents to them and to the societies and economies that they oversee. That is certainly not unique to the UK. For most countries, including the UK, the say-do gap, in that spending on AI, is largely confined to innovation budgets and stuff like that within Departments; it is not necessarily the mainstream activity of what they are doing.

One of the points that we need to address is not just what we invest in or how we do that in the long term, providing multi-annual funding, long-term certainty and so on, but the leadership to disinvest from certain things. Going back to the point about risk omission and commission, in most countries—this is certainly not picking on the UK—I do not see a willingness to say, "Not only are we going to invest in AI but, precisely because we are investing in AI, we are not going to invest in these three things over here. We believe that we can achieve the mission or deliver effect fundamentally differently, so we do not need those lumps of metal that we used to have, because we will have a fundamentally different approach." That is a lot harder, because we have so much sunk cost financially, bureaucratically and politically, in the stuff that we are already doing or have already started that it is very hard to walk away from that.

That is the real challenge. Bringing this back to defence, it is not just how can we bring in those future capabilities that we want, but what do we do



with the legacy capabilities and programmes that we already have? Some of those are going to need to endure—to be clear, some of those things will be there for another 20 or 30 years—but with some of them, we will have to slay some sacred cows, to get rid of them to make space in budget and so on for other ways of doing things.

To bring that back to the Ukrainian example, that is precisely what they have been doing when they started saying, “Precisely because we don’t have certain capabilities”—like an air force of the size and scale that they would like, so they have had to think fundamentally—“how can I achieve the same mission using different types of capabilities?”

Q16 Sarah Atherton: James, I want to pick up on something you said in answer to Derek’s question. You felt that procurement needed to sit with SMEs, sit with frontline command, to speed up and de-risk acquisition. Are you aware whether the MoD’s Commercial X is doing that?

James Black: There are lots of good examples—not just in strategic command and things like Commercial X, but across the frontline commands—of lots of more agile and innovative approaches to acquisition. That is true in other countries as well, and the UK can certainly learn from many of its allies and partners. The point, again, is to achieve scale.

We are very good at working on a small-scale, experimentation programme to develop, field and test some sort of autonomous system, for example—putting it out on Salisbury plain, flying it around, driving it around, whatever it is, and deriving some lessons from that—and doing that in close concert with industry and some SMEs. The challenge is then how we scale that. How do we bring that into our defence budget and our programming when so much is already taken up with buying the stuff that we committed to buy five, 10 or 15 years ago?

Again, it is about joining up the different efforts that are going on across the different services or different top-level budget holders. As has been mentioned, sometimes the Navy, the Army and the Air Force will all be looking at a similar problem, but approaching it separately.

A further layer of complexity is how we also plug in with cross-governmental efforts. As has been mentioned, in many cases AI on the operational and tactical level is quite bespoke to the military setting, but when we are talking about the back-office stuff—AI use in support functions, or enterprise and business management by the Department—a lot is not specific to the MoD; it is generic to any Department or any large organisation. Bringing in AI to improve HR management of civil servants is not defence-specific; that is a broader governmental thing.

Relatedly, there are going to be areas where you want to collaborate with allies. AUKUS has been one example, which I mentioned earlier, but also working through NATO or with European partners. The challenge is that we need to do all those things to join everything up to give you the scale, because if you do not have scale, then you are not going to do it at a level that makes enough of a financial impact to support businesses and does



not deliver enough of a capability to have an impact. Also, if we are achieving that scale through working with lots of different actors, that is itself complex.

There is a long history of multinational or cross-governmental procurement programmes where, because you have to align different departmental and national budgetary cycles, you have to get everybody on board and everybody's got subtly different requirements, which they all insist are desperately important. All those things add costs, complexity and time, which are anathema to what you are trying to achieve with delivering AI at pace. There are some difficult trades here and, despite good green shoots of good activity, we cannot say that that has been scaled yet.

Air Marshal Stringer: To find a positive in all this, because we were asked about cross-government and funding, where there is clear political will and an output that has to be driven to, we can do it. I would cite, positively, the Office for Security and Counter-Terrorism in the Home Office. It responded a few years back to that dreadful uptick in terrorism incidents, and suddenly all those Departments that used to corral and husband their own data had to share—across the intelligence agencies, the National Crime Agency, HMRC, the Passport Office, and so on. Suddenly, because there was an immediate threat and that immediate threat was persistent, and there was massive political will because of the costs of suffering dreadful terrorist incidents, we could do it.

It is worth looking at how some of those good models came together. I agree with James's point about how, when things work, often it is because little, heroic, self-taught pockets fight something through, or as I said, it is where you suddenly have a persistent but imminent threat and real political will to do something about it. I am not counselling here that everything is so broken that we cannot do things; I would say, look at the incentives and at the structures, and ask why we can do it in some areas, but do not seem to be able to move so quickly in others.

Chair: Before I hand over to Jeremy Quin, I ask everyone to be mindful of time. I encourage some brevity, because lots of Members will be leaving at about quarter to 1, and we still have quite a lot we need to get through.

Q17 **Sir Jeremy Quin:** In the light of that, my question is in three parts, although you have all tackled parts of this in answer to colleagues' questions, so if you feel there is nothing further to add, that is absolutely fine, but it gives you an opportunity to expand if you wish to do so. On the Air Marshal's comment about saying and doing, everyone has done a lot of saying: everyone recognises the incredible importance of AI and, I hope, the need for speed, and they have all been talking a good game—but there is very limited sweat across the MoD and the TLBs. Notwithstanding that, we have the Defence AI and Autonomy Unit, the Defence AI Centre and individual units across each of the TLBs.

Of my three related questions, initially to the Air Marshal, the first is, does this cohere? You have already given us your answer—it doesn't—but



HOUSE OF COMMONS

how do we improve on that? Is it better or more centralisation, or is there a risk that by pulling things together to the centre, we lose innovation and the warfighter being better in the loop? That is one question.

On the second question—it and the third probably come to Mr Black and Dr Soare—Dr Soare referred to 3,000 small businesses, which is a huge profusion, but none of them has unlimited resources. Many are fairly unresilient. When they look at this great mass of stuff going on inside the MoD—one of you talked about the demand signal; is there one, or is it just in the “too difficult” box? There are many easier customers than the MoD, with many less cumbersome processes. Does the sector of vast numbers of SMEs, nearly all with cash-flow issues, as you described, just say, “That's too hard,” or is there a way to ensure proper buy-in? If so, with whom? Who do they turn to? Is that demand signal clear, and is it clear who they should be speaking to?

There is a third related element. Mr Black referred to commission and omission. By way of anecdote, I remember personally saying, “Look, if we're doing exciting, innovative things, some of them must be going wrong. Why do I never hear anyone coming in and saying, “Sorry, we're not getting anywhere with this project. Let's close it down and recycle the cash?” It never, ever seems to happen. So what you get you hang on to, and you constantly pursue it, even if it is never going to work. That gives rise to a further question, which is the third element.

You refer to open architecture. You refer to preventing vendor lock-in. All of that requires a really intelligent customer. Do we have it? Those are the three elements.

Air Marshal Stringer: As you said, I have half answered this already. You asked whether, if you remove the warfighter, it would—

Sir Jeremy Quin: Do we lose something by having centralisation?

Air Marshal Stringer: Yes.

Sir Jeremy Quin: It seems obvious that we bring everything together, and there is one really great centre for AI, covering all the services. But would we lose something by doing so?

Air Marshal Stringer: No, because at the moment everybody wants to take an unmanned vehicle to Salisbury plain and put some stores on it and see it driving around, or a handheld drone on Salisbury plain. Well, that is great—a little tactical vignette. What we are really looking at here is asking big, at least operational-level, questions, like, “How do I get the Russian navy out of the Black sea when I don't have my own navy?” There are other people who can solve these things.

I wrote a paper just before I left, on a thing called the defence experimentation group. If you just look at pure experimentation, that is not what the frontline does on exercise; just £800 million for Dstl and £200 million in my own organisation every year—that is £1 billion spent on experimentation. It was a free good. No one is tracking it. You track travel and subsistence better than you track that. So I would centralise more,



HOUSE OF COMMONS

but I would centralise more around a warfighting headquarters that has some real operational, strategic-level problems to solve, such as the air and ballistic missile defence of the UK, and take those big problems, and with the SQEP that I do have, get it working together and then grow it from within.

At the moment you just have little pockets where people come in and then move out because they are posted away, and they have just learnt a bit about AI, and then they go back and do a personnel job or something else. You need a critical mass, and you need some big, operational-level challenges of how we will fight across five domains. Only then will you get to throw away the analogue Christmas tree and start building that core digital architecture—but more importantly, the ecosystem around it, with SMEs, co-ordinated by primes, working on big operational-level problems that will drive defence forward. At the moment it is not coherent enough to do that.

Sir Jeremy Quin: Do you have anything to add about who SMEs talk to and whether there is an intelligent customer when they do pick up the phone?

Air Marshal Stringer: I repeat my last point. There are one or two people around the place who have become pocket experts. They become expert in little bits and pieces, and then they get posted out after 18 months to two years, to go on to another, career-enhancing job. I ought to give the defence artificial intelligence centre the opportunity to at least grow itself, and at least we have done that. I would throw academia in here and—a personal hobby-horse—I would try to start the information warfare group in the defence academy, and actually have a nexus there around the defence cyber school, where you could have a continuing presence and grow a knowledge base. All those things are possible and are pretty much free.

Sir Jeremy Quin: Thank you, Ed. Simona or James, have you anything to add?

Dr Simona Soare: The SMEs and start-ups that we talk to in the UK and Europe are saying that defence is a very difficult customer. It is very difficult to understand. It is very difficult to understand the entry point. It is very difficult to understand what the MoD really wants, and what the problems are that it tries to solve. To give you an example, if you are an SME in the UK now, you are more likely to end up with a prime-driven or a private investor-driven accelerator. They will act as the gateway to defence, more so than doing it by yourself, simply because it is too difficult.

The opportunities to engage with MoD are rarer in comparison to the pace at which these intermediary accelerators are working. That is one of the big challenges—to reach the SMEs and also for the SMEs to understand what the mission request is from the MoD.

Sir Jeremy Quin: What you just said seems to be backed up by the DASA



statistics. Anything further, James?

James Black: In terms of the demand signal, I appreciate that it is a challenge in defence, compared to more hardware-centric parts of the departmental budget, to communicate your needs in a five or 10-year equipment plan. Clearly, we do not know, as we've talked about, what either the threat or the technology landscape will look like in such a fast-moving sector. I do think, none the less, that defence can get better. It has started to do this recently, publishing the use cases, tasks and missions that it wants AI to go after. The Air Marshal gave the example of the Ukrainians dealing with the Black sea fleet. That is the space we need to get into.

There is a recurring tendency in defence, applying to both hardware and software procurement, to over-specify what the solution looks like at the requirement-setting stage, and get down into technical specifications from people that really should not be touching them because they are wearing uniform—they are not on a factory floor somewhere in industry. That is a real danger because that doesn't give the space for new companies to offer innovative solutions to solve the problem in a different way. We are saying, "We need a lump of metal that's this big by this big and flies at this speed and has this many weapons." They might say, "Actually, there is a potential fundamentally different approach to delivering that," but they are crowded out because of the way we have written the evaluation criteria for the procurement. That links back to your point about being an intelligent customer. That is quite a tactical example of one of the challenges.

There is a more strategic challenge of being an intelligent customer. It is not just how do you ensure that you have the SQEP—you mentioned the technical skills and knowledge and so on—but how do you ally that with incentives, career management and cultural issues to enable people to take the risk to fail fast and fail early? Back to your point: we need to get to a point where people are actively rewarded in their career for trying something different for six months and then going, "All right, that didn't work, but it was worth doing. We've learned something from it. We've extracted that learning and fed it into another programme, so it's gone somewhere, and now I've shut it down," without seeing that as a disastrous management issue, which I think is how it would generally be seen nowadays.

Finally, there is a related point around understanding how people can influence and engage their suppliers, because we are moving beyond a traditional, transactional customer-supplier relationship to a much more enmeshed one, if you are working on things like software development and iterating over time, and we need to understand the levers, because this isn't working with BA Systems or Lockheed Martin; this is working with fundamentally different types of companies, which look to the MoD for very different things. The MoD has to be realistic about its level of influence and take the appropriate action to maximise that level of influence.



Q18 Mr Francois: Mr Black, your exposition a few minutes ago—about how incredibly risk-averse the MoD civil service has become and the really damaging effect that is having on our defences—was brilliant. I will not repeat it; it stands on the record. The culture within the Department itself is now a serious impediment to the defence of the realm. I hope that people will refer to what you said, because it was absolutely right.

Jim Mattis, a former US Defence Secretary, talked about doing things at the “speed of relevance.” It has been a very strong theme in the Committee this morning that when it comes to AI, we are not operating at anything like the speed of relevance; the bureaucracy in the MoD is incapable of keeping up with the speed of the development of this stuff. Is that a fair statement?

James Black: Yes, but with the caveat that we don’t have to, as the MoD, be as agile as the most agile Silicon Valley private sector firm; we need to be agile enough to deliver the defence of the realm against our adversaries. Although more authoritarian countries like Russia and China have certain advantages in terms of their ability to force things through their defence bureaucracy—clearly, people have slightly different approaches to things like promotion and career management at the top of the Kremlin—we need to be faster than the other guy. We need to recognise on the one hand that in Government, we will never be like the private sector, and there are good reasons for that: the public sector and defence in particular deliver fundamentally different outputs, and our attitude to risk in defence is always going to be different from that in some Departments. But absolutely it needs to go much faster.

Q19 Mr Francois: Forgive me, but a constant theme of this Committee—just look at our recent report, “Ready for War?”—is that defence often delivers outputs very badly. As you said, the Ukrainians are fighting a war of national survival, but it is not impossible that within some years we may have to do the same, so we should learn from them. Ed Stringer’s evidence was very powerful about how incredibly ponderous it all is. Within the MoD, we have a mechanism for doing things much faster, which you touched on. We used to call them “urgent operational requirements”, but now the jargon is “urgent capability requirements”. That is the same thing; it means that you define that you have got to have something quickly, you drop all the procurement bureaucracy—I nearly said something else—and you just get on and deliver it. Surely AI is an area where we should use UCRs much more frequently, shouldn’t we? If we have that accelerated way of procuring in the MoD—a methodology that people understand—surely AI is a classic area where we should apply it.

James Black: I agree. There are a bunch of different ways of doing that—UCRs is a good example—even within the existing rules. You can read the relevant Joint Service Publications, of about 600 pages online, which talk through all the procurement rules of the MoD, and there are quite a lot of different ways of being more flexible within that. It is not necessarily that different mechanisms or procedures are needed; it is people having the leadership, the culture and the attitude to risk in order to embrace the



HOUSE OF COMMONS

ones that are more focused on urgent delivery, but that goes back to the point about people feeling empowered to take those risks. In the Ukrainian case, they clearly are, because they have an urgent operational need. What we need to achieve as a Ministry of Defence is that same level of urgency pre-war rather than during a war. I would rather us get good at that now than get good at it in however many years' time when we find ourselves in our own existential war.

Q20 Mr Francois: To your point, the Defence Secretary said at Lancaster House that we have moved from a post-war to a pre-war age. That got pick-up around the world. When I was in Washington a couple of weeks ago, the Americans were using that phrase, but if you will the end, you must will the means. The official policy of the Government—of the head of the Defence Department—is that we are in a pre-war age. That should mean that all sorts of changes then follow, shouldn't it?

Dr Simona Soare: Historically, a pre-war age was measured at about 10 to 20 years, at least in my field of international relations. That, I think, goes back to the cultural issue within Government, which is the expectation that it is—I believe the integrated review calls it this—a once-in-a-generation modernisation that is expected to take that long. It is expected that our allies and potentially some of our adversaries might take that long to modernise their equipment. Whether they are actually doing so or not, the Air Marshal and James were talking earlier about the fact that they are moving significantly faster.

One thing that I would like to emphasise in relation to that is that there is also an element of how you consider not just our procurement, but how you consider the scalability of our projects. We are very much focusing our procurement on a vertical scaling solution—everything that we buy needs to be in these many numbers and these many lines of code—whereas, as I think was alluded to by James earlier, the key issue for artificial intelligence and data-driven technologies is to scale them horizontally. That is one of the one of the key challenges. The other key challenge that the Air Marshal alluded to was getting the services to work together and to agree on technological solutions that are automatically, by design, made to talk to each other and be compatible and interoperable.

Q21 Mr Francois: I will come to Ed in a moment. There is an old saying in the military, which he will recognise, that the enemy has a vote. The integrated review is far too ponderous. It has already been overtaken by events and is basically out of date because of what Russia is doing now. They have just murdered Navalny. So this very ponderous MoD document is completely overtaken by what's happening. It is not a document written at the "speed of relevance". I am sure the Americans had some long-term upgrade plans the day before Pearl Harbor, but they all went in the hopper the following morning.

Ed, in terms of trying to speed all this up, if you were king for a day, how would you do it?

Air Marshal Stringer: On the tech side, I would take what you said about urgent operational requirements. I ran that programme for two and a half



HOUSE OF COMMONS

years in a previous job. That was slanted at commercial off-the-shelf, and it cut through commercial rules of competition and other things. What I would say here is that you need to move at this speed, but you need to follow the minimum viable product method. What you actually need is to create a series of battle labs—I coined the phrase “drone farm”. What we are seeing in Ukraine at the moment is the ability to produce drones that are good enough. Now there’s a phrase—good enough should be good enough, and good enough is better than the opposition is today. If I go for the perfect tomorrow, it will probably be out of date. That is the way our system is structured now.

I think I’ve already covered this second point, so I will be very brief. Where the Americans and others win over us is in having combatant commanders—some of them have been referenced already—who have real-world problems to solve. At the moment, we structure defence around the inputs of the services themselves, who of course all have either a land, air or maritime-centric view of what the next war will look like. So I would pull those two things together if I was king for the day. I would have a stronger military strategic headquarters, setting real-world problems. As you say, it has a different worldview between 6 and 7 December 1941.

I would create ecosystems that would be different for the sort of tech. As we are seeing, we still need to produce old-school artillery shells, and the way you procure and produce those would be slightly different. But where a lot of your capability is software-driven, you need to work with industry minimum viable products. The battlefield is the battle lab, and you need to get into an iteration cycle that is quicker than the opposition. I will shut up on that because that is really just Darwin, isn’t it—the survival of the fittest.

Q22 Mr Francois: We have had excellent witnesses this morning. You have all said very important things, if I might be so bold. Ed, you talked about things being good enough. When the Committee looked at procurement, we found that one of the classic problems is that the MoD always goes for the exquisite solution, which takes so many years to come into service that it is often obsolescent by the time it arrives. You are saying that in the AI field we should do much more battle lab work, do something quickly, make sure it is good enough and fit for purpose—to use a military phrase—get it into service and then move on to the next evolution. Have I characterised your advice accurately and correctly?

Air Marshal Stringer: Yes, in the areas where it is appropriate to do so. The clever bit is how you mix and make the long-term stuff where you are going to have sunk costs. Once you have built the titanium shell of the submarine, you are stuck with that for a while. How, therefore, do you make it adaptable? How do you make it future-proof? We have already discussed certain things about open architectures and what-have-you. Where things can be software-driven and the hardware can probably be 3D-printed—which is why I am concentrating on drones in Ukraine at the moment—it is on the intellectual property that you need the MoD and the people planning the campaigns to be working with industry to work out what industry thinks it will be able to do tomorrow, and lead-turn that and



produce those things. The battlefield is the battle lab. The results come back and you iterate forward. That is a different model to writing a spec, then taking years down the CADMID cycle. So what you are doing is building a national capacity where at the end of it you have warfighting countries. Who else is on this Committee here? I suppose this is a version of the vaccine taskforce: you have a real problem to solve, you have to work with industry and you have to adapt—actually, even more so than the vaccine taskforce. Once drones, for example, are out there in a battlefield, they have a lifespan of only a few weeks before the opposition learns how to defeat them. Therefore, you have to keep iterating and keep ahead of the game. At the moment, we don't move quickly enough to play that game, let alone win it.

Mr Francois: The key thing about the vaccine taskforce was that from day one there was a sense of urgency—an imperative. I sometimes think the MoD civil service simply doesn't understand that any more.

Q23 **Chair:** Thank you very much, Mark. The panel will not be surprised to know that this Committee is also looking at GCAP, which has an in-service date of 2025. In the light of what you said about procurement, what are the implications for that programme?

James Black: Building on what the Air Marshal just said, I completely agree about MVP¹ and so on. It drives you towards a fundamentally different approach to what you are designing and how you conceptualise capability. You are no longer buying an aircraft and calling it your capability once you have stuck a pilot in it, got a runway and all these other things. You are instead thinking about how you can provide certain effects over the life of a programme, and how the hardware and software can be designed, configured, upgraded and refreshed over time so that the aircraft's capabilities in year one versus year 10 are radically different. You design the hardware to accommodate that, so you can stick modules in and out and change things over time relatively easily, and you design it much more around the software.

How you develop that obviously drives you towards thinking about a system of systems. You are thinking less of a platform-centric approach and more about disaggregating the effects that you want to achieve and the missions you want to undertake into the underlying systems and sub-systems that you need to get that. You need some sensors, some weapons, an energy source, propulsion and so on, and you start disaggregating all that.

That is obviously challenging to do from a programme governance management perspective, a commercial management perspective and so on. Linking back to the previous questions, there are potentially some really important strategic and economic benefits. The strategic benefit is that if we are moving to a model that is less about delivering a gold-plated solution and more about delivering something that grows and evolves over time—it is more about maintaining an industrial and governmental

¹ Minimum Viable Product



HOUSE OF COMMONS

capacity to innovate at pace. And that has a really important deterrent effect on our adversaries, because they know not just that we have aircraft on day one of the war, but that we have the underlying industrial base and intellectual horsepower to come up with something innovative in week six or month six of the war, and do what the Ukrainians are doing.

It also imposes intelligence and R&D costs on our adversaries because they don't know quite what target they are shooting. Currently they know, "Well, we need to blow up some Challenger tanks, some ships and whatever." They have a relatively fixed idea of the UK's capability. If we can show ourselves to be capable of much more agile delivery of capability at the time, including in things like GCAP, then they don't know what that capability will look like in year two of the war. They don't know how it will evolve, so they have to invest themselves in far more capabilities and areas of research to prepare for the directions in which we might take our future capabilities.

Briefly on the economic benefits—this is certainly central to GCAP—as soon as you move away from a platform-centric view and start thinking more about the lifecycle of the capability and how it can grow over time, you also introduce new opportunities for things like exports and partnerships with other countries, because they no longer necessarily have to partner on the whole programme; maybe they can partner on just one sub-system, or maybe you can export that sub-system to them and work with some other country's platform. Getting that right brings a lot of broader economic benefits and opportunities.

Chair: Dr Soare, are you able to add anything?

Dr Simona Soare: To follow up on that, this transition would very much benefit the ecosystem as a whole because it would strengthen the demand system across the supply chain and would allow particularly smaller companies in the supply chain to pitch into broader projects. The important thing for me is to understand that projects such as GCAP can capitalise a lot more on what we are seeing in terms of complementing governmental investment in key areas such as hardware, where it is very difficult for private capital to approach sustainable, with private capital investment in the private sector and in the software domain.

Artificial intelligence is being hyped, whether we like it or not and whether it is beneficial or not, for defence purposes—but it is beneficial because the private sector is interested, the investors are interested, in it. We are seeing a steady rise in investment in artificial intelligence. Across the Atlantic, it is doubling every few months. In the UK, it is doubling every couple of years. That stream of investment can be better leveraged by Government: "How much am I leveraging from private investment in the budget that I have allocated for my innovation, my project, my capability development?" To give you an example, the US have started to perfect this to the tune of about \$30 billion in terms of private capital investment in whatever programme worth at least \$1 billion they are putting forward through the Defense Innovation Unit. That is a huge portion of investment,



HOUSE OF COMMONS

of money, that they are leveraging. They don't have to spend it; it's already there.

By comparison, in the UK and across our European partner countries, we are seeing that proportion being significantly different. For every billion that the defence establishment in the UK and Europe are spending for defence development, capability development, they are leveraging about 300 million at best in private capital investment. Moreover, the return on investment when it comes to governmental investment in research and development for digital is about—very optimistically—0.5%, which is not sustainable over the long term.

So thinking about a software-driven or software-first approach to capability development, to system development, would also mean a different way of approaching how we budget for it—the procedures that we are putting in place for procurement—but also our engagement with the AI ecosystem and how that AI ecosystem is able to leverage both commercial and governmental financing solutions.

Chair: Edward, do you want to add anything?

Air Marshal Stringer: I know you want us to be brief. I have two points. To reinforce that last one, it is indeed some of the SMEs who have come to me and said, "The MoD just isn't leveraging the capital I could raise on the markets. If it can partner with us and I can go out to the markets and say, 'I'm working with the MoD on this,' for every pound the MoD puts in, I'll raise another five." That was introduced at high level a couple of years back. It went nowhere, and these companies are frustrated. In other words, we put a pound in and we get a pound's worth back; it's very transactional.

The second point is just to reinforce how wars are actually won. I wrote a thing called the case for force development, because it is force development that actually wins wars. Wars are not won by the few bits of bespoke, maybe exquisite equipment that you start the war with on day one. Those are useful for what we might call special military operations of short duration. I do think we need to flip our mindset on this and set the MoD up as essentially an organisation built around problem solving and force development that can then move quickly, and the move to the information age should be the strategic shock, the impetus, that forces us to do that.

Chair: Thank you all for that.

Q24 **Sarah Atherton:** I was going to ask whether the MoD is AI-ready, but I think we have already answered that, so, in the interests of brevity, I will just pick up, Ed, on something that you mentioned way before. There is no single owner of AI in defence; the overall management, I understand, is shared between the Defence AI and Autonomy Unit and the Defence AI Centre. Do you think that, given everything that has been said today, we need a named person, a senior responsible owner, in defence to take this on?



Air Marshal Stringer: I don't know. I have thought about that, but the trouble is that it can then become very abstract. We had this a few years back when I ended up being SRO for drones, when I was on the air staff, but drones don't exist in a vacuum unless you're thinking of toys. What you are actually talking about is the imaginative use of autonomous systems in a whole load of areas, so drones should be a subset, or thinking about drones should be a subset, of those who really understand, for example, anti-submarine warfare. On having an SRO for AI, that should really be the Chief of the Defence Staff. They should ensure that all three services are building to a common vision and will come together in a way that creates a warfighting ecosystem that merges with the physical, as we see in Ukraine in the huge amount of heavy kinetic warfare going on. However, the connecting tissue that links them together is the bit that amplifies that, makes it work and makes it work better than the opposition. If you have an SRO just for AI per se, there would be a lot of conferences and talk about the abstract, but you would not actually be driving to get things to work.

When I talk to people in industry—this is my final point—they will always work with the customer and say, "What is it you are actually trying to achieve? Let's find a practical solution," rather than starting to talk in generic terms about AI in the abstract and trying to sell them some AI. How can they? It must be involved in some form of problem solving. I would look at the structures of how we measure defence outputs first before I went down a route of having someone to control the inputs.

Q25 **Sarah Atherton:** I understand the Chief of the Air Staff is the senior innovation officer for defence. Is that something that could be added on to that role, or does it need to stand alone because of the breadth of the subject matter?

Air Marshal Stringer: I am sorry—I wasn't aware that he was. I am not sure how he could be, because that must be a head office function.

Q26 **Sarah Atherton:** We learned that the other day, the Chair and I, and we did not know that either. Thank you, Edward. Is there anything else you want to mention that has not been mentioned already about MoD AI-readiness? How could it improve?

Dr Simona Soare: One thing to highlight is the fact that the defence AI strategy says that within 10 years we will end up with an AI-ready force. When we get to 10 years, if we are only AI-ready, we will have been so far behind the eight ball.

James Black: One of the really important things will be implementing a lot of the findings of the Haythornthwaite review and the broader changes to people management in defence. We keep coming back to the issue of things such as SQEP² and culture, and that really reflects the fact that we often talk about AI as if it is a technical system. It is not: it is a sociotechnical system. There is a key human factor in terms of both the individual user or team of users and all the people who sit behind them in

² Suitably Qualified and Experienced Person



HOUSE OF COMMONS

head office buying the thing in the first place, hiring the people, paying their payslips and all those sorts of things.

There are some huge opportunities from AI in terms of how we drive improvements in the wider management of the Department. People will have different requirements. Clearly, the military requirement for AI at the tactical end is very different from the HR team's requirement. However, it requires people to have at least the basic AI literacy to understand enough about the technology to understand how it may be disruptive or beneficial to their day-to-day lives and the work that they do. They will then be able to do their part in changing the culture as a senior leader, a procurement person buying the thing, or an end user coming up with clever suggestions for how that could be improved, which you can feed back into the iterative process of software development. That is a broader challenge around tech skills and STEM skills in defence in general, or in Government or the UK in general.

Lots of positive initiatives are under way but, going back to the point we were making earlier about urgency, being willing to go far enough with that will be really important. We can have the best structural and procedural ways of incorporating AI as an organisation, but if you don't have people who are able to navigate and use them they will default back to poor practices.

Q27 Martin Docherty-Hughes: I wonder whether we can talk about infrastructure briefly, given the time constraints. Do you agree that the UK has the infrastructure necessary to support the development and deployment of advanced AI—for example, computing power and access to appropriate datasets on which programmes can be trained? James, you first.

James Black: Clearly, this has been a big emphasis for the UK Government over the last six to nine months more generally. Obviously, there have been announcements about investing in more compute domestically; crucially, it is about who has access to that compute. If you are in the academic sector, for example, it is obviously harder to get access to that than if you work in a big tech company with billions and billions in revenue and have a lot of that stuff in-house. That has obviously been positive; clearly, more is preferable.

In terms of data, obviously there are some specific challenges with defence around what you can share. Some of those are sensible, security-related concerns, but many are more bureaucratic or data management concerns—if the data is not cleaned, interoperable with other datasets or sharable in a useful, timely format. We are also seeing the rise of synthetic data, which means you can do away with needing real-world examples for certain things, which can be beneficial. There have been some positive steps. We have already talked a bit, particularly with the Air Marshal, about the data strategy. Defence is moving in the right direction in terms of its broader approach to data, but is it in the place any of us would want it to be in now? I think the answer would be no, and I don't think that would be controversial to pretty much anyone in the



HOUSE OF COMMONS

Department. That is a challenge that many large organisations face, not even just in the public sector. Clearly, you need that infrastructure.

Relatedly, there is a question about how you manage data and infrastructure as the key determinants of your sovereignty and security of supply. We are used to thinking about things like security of supply in the traditional defence industry: do I have the factory that makes the artillery rounds—yes or no? It becomes harder when you start to talk about slightly more ethereal things like software development and compute, which has a hardware element but is often shared across borders. A lot of the value is now wrapped up in data rather than something more tangible, which requires a slightly different mindset and approach to things like our industrial policy, and, crucially, to how we work with allies and partners to ensure that the UK gets a good return from those partnerships with other countries. What we don't want is to get into a position where we have embraced AI, which is great, but we have made ourselves incredibly dependent on some other countries, which may or may not be reliable providers in future.

Martin Docherty-Hughes: Simona, do you want to add to that?

Dr Simona Soare: Yes. There is an element of catch-up in terms of ensuring the infrastructure element of software and digital transformation in the UK. If we look at international metrics and comparisons with other countries, definitely the area of infrastructure and semiconductors has been a weakness in the UK AI ecosystem. Recent efforts by the Government are targeting that, and they are aware of the gap and the weakness there.

Another element to highlight relates to the structure of the ecosystem itself. There are significant numbers of AI companies that work in the infrastructure and associated AI infrastructure domain. It is about 30% to 40% of the companies, so it is not a small proportion overall. However, they are very dependent on external actors for key areas that they need. When we talk about compute and semiconductors, these are rare earths and critical mineral-heavy elements for which the UK depends significantly on other, potentially adversarial, countries. Those elements of dependency and gaps within the ecosystem and the critical supply chains are incredibly important when it comes to us understanding the health of the infrastructure that supports the digital and AI ecosystem in the UK.

Martin Docherty-Hughes: Edward, do you have anything to add?

Air Marshal Stringer: Dare I use the term "industrial policy"? It strikes me that if within all sorts of national strategies we are pushing the tech industry, because the tech industry underpins life sciences and all sorts of other things, we would be making it very easy to open huge amounts of large cloud storage in this country and make all sorts of things possible. I absolutely agree with what was said previously about this being a socio-tech problem. In fact, I think the first Development, Concepts and Doctrine Centre paper on human machine teaming was written in 2017.



HOUSE OF COMMONS

On data, it tends to belong to someone, and there is one think that the Government could do. I am branching a bit away from defence, but so much of this stuff is dual use, so you cannot hive it off. If one thinks about the NHS, just think of the data that we have, almost uniquely, because of the size of the NHS, and that is just one example. The conversation gets very difficult as it gets politically challenging. People ask, "Who owns my data? Who will be able to see it? That's all very private." So the one thing the Government could do is really start that conversation. It is easier in Estonia, when there are only 1.2 million of them and they have just had, in 2007, a huge cyber-attack from a neighbour that used to own them—you can cut through some of that. It sounds so easy to say, and I know it is difficult to do, but I think these are the areas where Government could set some conditions for future success.

- Q28 **Martin Docherty-Hughes:** Can I link the issues about industrial policy and computing power? You will forgive my ignorance, and any of you may correct me, but if you are looking to maintain computing power, then we need a sustainable approach to energy production, and one that is safe. In terms of that being part of the infrastructure to maintain computing power, if there is a data centre owned by Amazon in Ireland, for example, which takes 11% of Ireland's energy production, where does that leave the UK?

James Black: Yes, clearly energy costs are a big driver for those maintaining data centres. There are some innovative ways in which people are trying to reduce the energy demand in the first place from those centres, including pretty out-there things around cooling, such as looking at putting them underwater so they naturally cool in the environment.

Martin Docherty-Hughes: Which we see in Norway, because they are getting 24/7 green energy for data centres.

James Black: Yes, exactly.

- Q29 **Martin Docherty-Hughes:** So it is not exactly new. Why are we so far behind when it comes to maintaining that energy structure?

James Black: That is a broader question on energy policy that is probably slightly beyond me, but I think it is absolutely a big thing. This goes back to the point about things such as energy, data and computing being the basic, fundamental building blocks of the current digital economy and of the future digital economy that we want to have—and, in turn, of defence. That is forcing us, in the purely military mindset, to think quite differently about security of supply and things like that. We are no longer just thinking about dependence around, "Where do I get my 155 mm munition?" Although clearly that is really important right now, we are also starting to think about, "Where am I getting my semiconductors?", and, as you have mentioned, "Where am I getting the rare earth elements?" Then, crucially, "What role does the UK have in that value chain?" It is equally unrealistic for an autarkical approach, where the UK tries to replicate all that itself, so how can we de-risk our reliance on global supply chains? That is where you get into the question of friendshoring—working with our



allies and partners, having mutual dependence with people we can broadly trust, and trying to de-risk that way.

Dr Simona Soare: But the current supply chains are actually very concentrated—and not with our allies and partners—and that is the major weakness of the ecosystem that is actually lying very hidden at this point. There was a recent Department of Defence-sponsored study in the US—done by Govini, if I am not mistaken—and similar studies have been done by the European Commission, across the pond in the EU, that have looked at the health of the supply chain when it comes to insuring or de-risking some of these dependencies. I am not aware of whether something similar has been achieved in the UK, but that would be very useful to understand where the single points of failure are.

In the US, they have determined that potentially about 80% of their secondary and tertiary levels in their supply chains are actually infiltrated by China and Chinese actors, which means that a smart adversary who knows where to put pressure can actually incapacitate that supply chain for defence pretty easily. It is the same thing for the European Union: they discovered that they had 100% dependencies on China for key elements of the critical supply chains for defence. Again, for a smart adversary that knows where to apply pressure, that is a pretty powerful structural element to play on.

Q30 **Martin Docherty-Hughes:** Thank you. Ed, do you want to add any final comments?

Air Marshal Stringer: I think that we can tie it very quickly back to AI here. As we have said all the way through, AI requires data—machine learning requires you to join data. You mentioned 11% of Ireland's energy; energy security is a vital part of national security anyway, and you've now got it for this. Because they underpin a lot of our defence, the data farms then become critical national infrastructure that needs protecting, both virtually and physically, so there are military and security tasks in that.

So yes, all these things are fundamentally and intimately linked, and it would be good to see some of that being addressed. I know that is outside of the scope of this immediate Committee, but I am not seeing too much yet—other than some eye-catching electric scooters and things like that—about quite how the MoD is going to power its military in even 10 years' time, when some other Government policies will have been met.

Q31 **Richard Drax:** Good afternoon to the three of you. I'm afraid brevity is needed because time is ticking by. There are two parts to my question. One is that our evidence tells us that AI skills and the talent gap in defence are an obstacle in developing a defence AI sector in the UK—can you comment on that? The second part is how do we or the MoD motivate AI professionals to work in defence when there is so much more commercial benefit outside in the private sector? Edward, perhaps you could give a brief response to those two points first.



HOUSE OF COMMONS

Air Marshal Stringer: The choice is thin nationally. I think I mentioned earlier that a very big IT company set up a big institute with one of our universities and could not find someone who is not Chinese to come and do a sponsored PhD in a national security-sensitive area—that took months to resolve. When you go down the chain to find Brits you can sponsor, you are at undergraduate level, so the choice is thin. GCHQ has exactly the same problem, and has for a long time.

I am more sanguine about working with defence. I have an idea involving reserves and sponsored reserves, because the Government can do things with IT that only Governments can do, and they are absolutely fascinating. A lot of this stuff is dual use—as we talked about earlier, if you get your force development, or spiral development, or the battle lab right, then these people can be an industry that provides national security and defence output. They would probably find that a very interesting part of a portfolio career in IT. So there are ways around this, as long as we can break up some of the legacy thinking on how military careers have to be structured.

Q32 **Richard Drax:** What about pay?

Air Marshal Stringer: If they are working for a company then the company pays what it needs to pay to keep them doing what they are doing.

Richard Drax: The MoD is never going to match the private sector for pay, is it?

Air Marshal Stringer: That is what I am saying. GCHQ has got exactly this problem with people who are very capable in the cyber sphere. Part of it is that people will do a bit of time with you, because it is absolutely fascinating—like how people go and work in the Treasury as youngsters and then go out into banking and so on; it is a fantastic grounding.

You need to think through how careers are managed and exactly what the offer is, rather than trying to recruit someone as a 16-year-old and get them to stay for 20 years doing nothing but sitting on a military pay while becoming a global expert in AI. That is not going to work, but as we have just said, defence outputs will come from working much closer with these start-ups on problem solving, and the start-up will pay the salary that the industry does. Of course, the MoD will end up paying for it in the end, but it is doing it through capability rather than through the salary route.

Q33 **Richard Drax:** Can you clarify the point you made about the Chinese? I found it quite shocking. Did you say that all you could find was Chinese students? Was that over two years?

Air Marshal Stringer: That's been true for a very long time. You look across at the university sector and there is a paucity of UK students coming through in AI. I shall defer to my fellow panellist from the university sector on this, but when I talk to the industry this is something that repeatedly comes up. I am aware of a couple of incidents where there has been a conflict because the university put forward someone who



HOUSE OF COMMONS

China—a competitor state—wants. But there are lots of Chinese students in our universities and they are a very large fraction of the student population.

Richard Drax: Dr Simona, perhaps you can pick up on that point.

Dr Simona Soare: There is a significant cohort of foreign nationals, particularly Chinese nationals, who are involved in STEM disciplines—and certainly AI. We have been in situations working on projects for defence where we had to make choices between the kind of talent we were choosing, depending on areas of sensitivity, and one was the origin of the talent itself. We had to choose a different nationality, even though the prominent candidate or the best in the field would have been a Chinese scholar.

Q34 **Richard Drax:** So we need to put a lot more emphasis on STEM in this country to get our people to take it up—is that what you are saying?

Dr Simona Soare: Yes, I think so. STEM needs to be made more attractive to UK nationals and our partners across Europe and the United States. There is a lot of mobility, as James was saying earlier, between the two sides of the Atlantic, so that can be capitalised on, as with our partners in the Indo-Pacific. That also has to be emphasised because they increasingly invest in STEM disciplines as well, and particularly AI. To me, it is also important to think not just along the educational pathway; how many AI skills can I produce within a year based on what the industry needs? That is a key element, especially when we are thinking about scalability within the ecosystem. One of the challenges we are hearing from companies, small and big, is the fact that if they need to scale up the production of anything in the digital domain, they will see significant shortfalls in the skills available to them. That sometimes leads to delays in them meeting deadlines with the MoD.

There are signals within defence that this is being thought of not as a defence problem but as a whole-of-society problem. That goes to what the Air Marshal was referencing earlier: the idea that you are not going to bring in an AI specialist and keep them on for 20 or 30 years and they will retire from the MoD or from the services. Rather, bring them on for a period of four, five or six years and then they will go back into the commercial sector. The whole of society still capitalises on that because the AI skill remains and is leveraged across different sectors. I think that is one key element to emphasise. There is also some movement on thinking in terms of how we bring in AI skills laterally when we need to complement AI skills in defence. Obviously, that is moving a lot slower than we would like, but at least there are people thinking about that.

There is one last thing that I would mention about skills, and I would put it into the weakness or lack-of-clarity category. I am not really sure, and maybe my co-panellists are better informed about this, but when we look at the data that is being put forward by the MoD, it is very unclear as to the level of skills we are talking about that we need to achieve, especially from the educational sector. Industry will tell us, “We need 25,000 people



HOUSE OF COMMONS

within a year. Can you deliver?" When it comes to the MoD, the defence ecosystem is a lot more closed off and we do not really understand. Does it need to be 3% of the force? Does it need to be 20% of the force? What is that level that we need to achieve, and will we know how to deliver it on time?

Q35 **Richard Drax:** So a flexible career structure is, I think, one of the things you are saying—to look at a very different structure of career to meet the demand. James, do you have anything to add?

James Black: I agree with what has been said. You can either increase the total pool of people from which you can draw, or you can increase the portion of that which you can access. The increase in the total pool is a national endeavour. Clearly, it is an education policy issue and so on, so it is beyond the MoD. But it is akin to what the Americans had in the 1950s and 1960s with Sputnik, where Eisenhower and Kennedy said not just, "We're going to go to the moon," but also, "We're going to all do math. Everybody at high school is suddenly going to do science and maths education. We are going to pump out far more engineers and we're going to set up a bunch of engineering departments at universities," and all those sorts of things, to drive growth in the total number of people doing STEM. That is the sort of thing that we similarly need, not just for AI but for other purposes.

Of course, the follow-on from that is how much of that you can access to bring it into defence. As has been mentioned, there are always going to be some security and nationality caveats that apply to defence and do not apply elsewhere, and we need to find a balance between attracting top talent to the UK—because that is a big driver of our AI sector—versus making sure that we have sufficient British-only AI talent to feed some of our unique needs.

Then it comes back to incentivisation. We have talked about career structures and I agree with all that has been said. On the incentives, clearly you cannot win financially compared with the private sector if you are the MoD but, as the Air Marshal said, you absolutely can provide some non-financial drivers: sense of mission, sense of purpose and getting to do some cool stuff, which matters if you are a geek, frankly. We need to make sure there are all the other basic enablers there as well: the housing offer, the support for families—all the wider things that go into providing a compelling employment proposition to people in defence. Most of that is not specific to defence; it is a more chronic challenge for defence, at the moment anyway.

But yes, absolutely: if we do not get this right, then you can have all the procedural and technical solutions you like, but if you do not have the people, we are not going to get very far.

Q36 **Chair:** I have a quick final question on AUKUS pillar 2. How can the MoD ensure that pillar 2 of AUKUS enables UK developers to collaborate with our allies, and how can we overcome US restrictions such as ITAR and not releasable to foreign nationals? Again, can we have quick answers, if



possible? I know it is a bit of a long question.

James Black: Myself and some Australian and US-based RAND colleagues have just produced a report precisely on this topic for the three countries, looking at how you can drive collaboration. The answer is, unsurprisingly, that you absolutely need a holistic approach. Particularly in the US, they have taken some really positive steps with the latest NDAA,³ moving to designate the UK and Australia as domestic to their defence industry, which unlocks a whole load of co-operation and sharing. They have taken steps in terms of looking at ITAR,⁴ which is the export legislation. Reciprocally, Australia in particular has been looking at shoring up its own national regulatory environment, particularly its security around things like tech leakage and foreign direct investment, which I suppose is a prerequisite for buying the trust of the Americans, to be willing to share more in the first place.

I think it is important that the UK continues to show vigilance in its own approach to those sorts of topics, and how we scrutinise things like foreign direct investment and make sure that the technology and IP are not leaking out of the UK, particularly where they have been shared with allies and partners.

You then need to align yourself conceptually and doctrinally with what it is you are trying to achieve with AI or autonomous systems—that is, what missions we are going after and what sort of effects and systems we are trying to deliver.

There is then all the usual stuff around building opportunities for experimentation, standardisation and interoperability, which we have touched on a lot already. That is things like having data standards, or at least interoperable data sets, and making use of some of the opportunities that are there in terms of things like testing infrastructure. Australia is a big place with not very much in a lot of it: it is great for test ranges, for examples, for different things. Obviously, it is then about providing a joint demand signal from the three countries—to the extent possible—where we agree on common requirements, rather than all going off with our own bespoke requirement. We can therefore procure from industry at scale and get the economies of scale that are obviously beneficial to the taxpayer in driving down costs, but also to industry in driving up revenues and bringing in private capital investment and so on. It is that whole suite from high-level legislative and policy changes through to some of the regulatory changes, down to quite tactical and practical things that can be done in the DoDs and MoDs of those respective countries.

My final quick comment is that AUKUS is therefore, between pillar 1 and pillar 2, an international-level transformation akin to what the Air Marshal has been talking about in the domestic setting of trying to move towards force development, spiral development, problem solving, agile,

³ National Defence Authorisation Act

⁴International Traffic in Arms Regulations



HOUSE OF COMMONS

innovation—whatever jargon you want to throw at it. AUKUS is trying to do that with allies and partners. Obviously, it makes sense for it to be those three countries currently, and they are aligned on a whole load of issues. With pillar 2, you have an opportunity to also bring in other countries on a project-by-project basis where they have interesting technology, and that is where you get into discussions with others such as Japan, which obviously has something to offer as well. It is a huge opportunity, and it can drive a different approach internationally but also domestically.

Chair: Dr Soare or Edward, do you want to add anything very quickly to that? No. Okay. Thank you very much to all our panellists, my Committee colleagues and all the Committee staff. It was really fascinating, and it is the first time the Committee has looked at AI, so you have given us a lot to think about and chew over for future sessions.