



## AI in Weapons Systems Committee

### Corrected oral evidence: Artificial intelligence in weapons systems

Thursday 7 September 2023

10 am

[Watch the meeting](#)

Members present: Lord Lisvane (The Chair); Lord Browne of Ladyton; Lord Clement-Jones; Lord Bishop of Coventry; Baroness Doocey; Lord Fairfax of Cameron; Lord Grocott; Lord Hamilton of Epsom; Baroness Hodgson of Abinger; Lord Houghton of Richmond; Lord Mitchell; Lord Sarfraz; Lord Triesman.

Evidence Session No. 13

Heard in Public

Questions 165 - 190

### Witnesses

**I:** James Cartlidge MP, Minister for Defence Procurement, Ministry of Defence; Lieutenant General Tom Copinger-Symes CBE, Deputy Commander, UK Strategic Command, Ministry of Defence; Paul Lincoln, Second Permanent Secretary, Ministry of Defence.

### USE OF THE TRANSCRIPT

1. This is a corrected transcript of evidence taken in public and webcast on [www.parliamentlive.tv](http://www.parliamentlive.tv).
2. Any public use of, or reference to, the contents should make clear that neither Members nor witnesses have had the opportunity to correct the record. If in doubt as to the propriety of using the transcript, please contact the Clerk of the Committee.

## Examination of witnesses

James Cartlidge MP, Lieutenant General Tom Copinger-Symes and Paul Lincoln.

Q165 **The Chair:** Welcome, Minister, Mr Lincoln and Lieutenant General Copinger-Symes. I am sorry that we have a rather long-range exchange this morning. In this room, there is no alternative, but I think the voice-raising system will work pretty well. You know the form for these sessions; they are being broadcast, they will be transcribed, and you will see a transcript afterwards.

Minister, thank you very much for your letter, which we may take as your opening statement, so to speak. May I thank you too for your officials' assistance at the outset of our inquiry and the written evidence that was submitted when our inquiry began? It was extremely helpful. Thank you. Could you perhaps start by introducing your colleagues to us and explaining how their roles fit into the subject of our inquiry?

**James Cartlidge:** I am the Minister of State for Defence Procurement. I have to my right Lieutenant General Copinger-Symes.

**Lieutenant General Tom Copinger-Symes:** I am the Deputy Commander of the UK Strategic Command, which contains many, not all, of the key elements of defence's approach to artificial intelligence.

**Paul Lincoln:** I am the Second Permanent Secretary in the Ministry of Defence and I have overall responsibility to Ministers on the implementation of the defence artificial intelligence strategy.

Q166 **The Chair:** In answer to our questions, do come in, both of you—under the guidance of the Minister, obviously. Do not wait to be asked.

A key area of government policy will be how AI and AWS systems are procured effectively, safely and with possible value for money. Your department has had a bit of a run-in, has it not, with the House of Commons Defence Committee? Perhaps I should declare an interest, as I was clerk of that committee for six years in the 1980s and its criticisms, I fear, have an extraordinarily familiar ring to them. It said that the present system was "highly bureaucratic, overly stratified, far too ponderous, with an inconsistent approach to safety" and "very poor accountability". Fighting back the reaction, "Don't hold back. Tell it how it is", how can you convince us that you are in the best possible position to procure AI and AWS systems that meet the sorts of criteria that we and you would expect?

**James Cartlidge:** Thank you, Lord Chair. It is a very good question to start with. As Minister of State for Defence, I agree that procurement is a very live issue. I am not sure that is necessarily different from the 1980s. I note your point about the committee. I should imagine that many of the questions from then are relatively common, and I think that successive Defence Ministers—there are at least two in the room—will have experienced similar questioning to that which I received in my two sessions to date with the Commons committee on procurement: the first was on aviation, and the second was on acquisition reform.

The key point to stress is that, on the one hand, there have been major projects where there have been well-known, well-documented issues. Probably the most notable has been to do with the Ajax platform, although in my first Statement I was very pleased to be able to confirm that that is now with the Army and that it is training on it in standard field training on Salisbury Plain. We recognise, however, that there have been many challenges, which is why had successive inquiries, most recently the Sheldon review, and we have been transparent about taking on board all the recommendations in whole or in principle.

If you step away from individual cases from which we seek to learn lessons, there are many positives. First, on the key point about the time it takes to procure, in the two years to December 2022 the average time for a project reduced by a full year. That is a very significant reduction in time compared with average procurement times. Also, fast forward to where we are now, with the very real military situation that we face with Ukraine. As you know, DE&S, which is based in Abbey Wood, is critical to the procurement effort, and I believe that under the leadership of Andy Start, the CEO, it has made significant improvements in productivity and in the way it operates. So it is seeing significant improvements. Probably the most important improvement was the speed with which it procured urgent requirements into theatre in Ukraine, which has been crucial in enabling Ukraine to maintain its fight, which we all support and will continue to support.

However, you are right: there is this overall question about acquisition reform. I am passionate about it. I was clear to the Commons committee that I recognise the need for reform. Whether it is AI or any other platform or area of procurement—I accept there will be certain factors in digital matters that Paul Lincoln may want to comment on—whatever it is, the most important principle is this: we have to be swifter and more agile, because otherwise we will lose the competitive edge against our adversaries. That is the single most important reason why we need to improve acquisition, and I assure you that we are very conscious of that.

**Q167 The Chair:** Before Mr Lincoln comes in, I wonder if you could both incorporate in your response the issue of shortening procurement times, because the development of AI is moving at a scary pace, and if you were procuring a main battle tank, a fast jet or something of that sort, you would be looking at decades from go to whoa, but you may be looking at months now. Could you give us a feel for how you will adapt to that possibly very demanding fast procurement turnaround?

**James Cartlidge:** On the general point, what you will hear mentioned a lot—it was in the Command Paper refresh and is something that I was clear on in front of the Commons Committee—is that rather than going for 100% what we call exquisite platforms from the off, we go for, say, 70% and then have spiral development. By doing that, you can have a platform—if it is technology—in service more quickly and then you can keep up with technical change. In the digital space, that principle of procurement will give us the greatest ability to effect the sort of change you are talking about, which is how you keep up with what you rightly

described as the rapid pace of change. We have also some very specific changes in respect of how we operate digital procurement, and as Paul Lincoln has a strong role in that, I will bring him in at this point.

**Paul Lincoln:** In the Defence Command Paper refresh, there is also a commitment to timeframes. So although there may be exceptions, it is stated in the Command Paper that as a department we would look to having no more than five years for conventional platform-type technology, and when it comes to software we would set ourselves a timeframe of three years.

Recognising that to be able to do that we need a more agile process, which, when it comes to software, will inevitably involve more small and medium-sized enterprises, we have put new commercial processes in place to support that. Previously, our commercial processes would probably take around 240 days to get something to contract. The new Commercial X process that we have adopted now takes us to an average of 120 days, so we have reduced the time by 50%, and our ambition is to reduce that significantly further, recognising the point you have made.

Similarly, given that there is a significant amount associated with small and medium-sized enterprises when it comes to artificial intelligence, it is not just about big primes. The department also has an SME action plan, which gets us into the position of engaging directly but also expects that our prime contractors will similarly engage with small and medium-sized enterprises that will support this.

**The Chair:** It is encouraging to hear that. In that process, is there likely to be a tipping of the original balance, so that when you are dealing with suppliers—obviously this puts a strain on your resources of assessment and contracts and so on—you are not so much saying, “This is the bit of kit we want. Can you provide it?”, which is obviously the foot the boot is on at the moment, but turning it around where an independent developer comes to you and says, “We can do this. Do you want it?”? How are you going to accommodate that in evaluation?

**Paul Lincoln:** We have a balance of different approaches that we can take. We can set requirements to industry. Of course we receive proposals from industry separately, but some of the work that we have done will allow people to say, “Is there a set of problems that we would like to discuss with people?” rather than saying, “Here is a set of specific requirements”. We can cover all three of those approaches under our blend of different commercial approaches.

**James Cartlidge:** If I may just follow up, because this is such an important point, obviously there is a lot of instinctive entrepreneurial creativity happening out there. Your point, Chair, is that rather than it just being, “This is what we want”—prescriptive—it is about how we encourage and incubate that sense of innovation, particularly with companies that may not necessarily consider defence procurement in their typical corporate goals.

The sort of activity that we have undertaken is in building what we call an ecosystem, where we have strong engagement with those sorts of businesses—SMEs. I ran an SME before I came into Parliament, so this is something that I feel passionately about. Take DSTL, which you will be familiar with. It is very important in AI and has organised the AI Fest, which has been increasingly very well attended by precisely the sorts of SMEs that we want to encourage to start thinking about defence applications. Ultimately, as you rightly say, this is an area of what you might call not just rapid change but disruption. We need to harness that, and the best way to do that is by engaging closely with SMEs.

I was in Poland on Tuesday for its main defence fair, its equivalent of DSEI, which is being held in London next week. It was so uplifting for me to see UK SMEs out there with fantastic technical offers at the cutting edge. We want to encourage those businesses to grow and to be confident of being able to procure through the MoD.

**The Chair:** As partners, you will depend hugely on highly capable personnel, which is a question that Baroness Doocey wants to follow up.

Q168 **Baroness Doocey:** I am interested in how the salaries that you pay your staff working on AI compare with private sector salaries. How successful have you been in both attracting and retaining key staff, the brightest and the best? It would be very helpful if you could give a couple of examples.

**James Cartlidge:** It is an excellent question. I think we are all aware that across the board there is a profound challenge in the labour market in finding good people—and, probably most importantly, in retaining good people—and it is not just about pay. In the Ministry of Defence, we are very proud of the fact that we have a very clear mission, but, ultimately, no matter what steps we take, we will never compete with the private sector potential that a person could earn. That has been a long-standing fact. People who work in the public service generally have other motivations, and we think that people in the MoD still want to work in defence, obviously for the services themselves, but because they are also motivated by patriotism and public service. However, I will turn to Paul Lincoln to talk about some specific steps that we have taken on AI to boost recruitment and retention.

**Paul Lincoln:** We recognise the challenge. As set out in the national government AI strategy, there is a set of questions on skills. There is a national skills shortage in this area, and the Prime Minister chairs the National Science and Technology Council, which looks at that in the round.

From our perspective, we fit into wider government frameworks on this, but there is a new data and digital framework, which provides for additional salaries for people who are under that framework. It provides for approximately an extra 10% on peoples' salaries. As the Minister will say, that does not compete with the sorts of salaries that people might be able to get in the private sector, but we recognise that that is the case

and that we need to do our best in upskilling not just individuals but the workforce as a whole.

I might ask Lieutenant General Tom Copinger-Symes to talk in a moment about some of the work that we are doing on digital skills for defence, but we are increasingly investing in how we can access private sector skills training. At the London AI summit, I announced that we had put a contract in place with Google, for example, to access all the training material that Google uses with its teams when teaching them data and digital skills. We also recognise that within defence, we want to turn digital and data skills into a profession, and we are establishing an AI skills head of profession role, which will sit within strategic command to help to make sure that we are setting the right framework across the department.

**Baroness Doocey:** I understand entirely that you will not always be able to compete with the major private firms and that you now have a system that will add 10%. Unfortunately, that does not mean anything to me. Can you tell me what percentage of salaries you are talking about with the 10%? Will that apply 50% of what the private sector pays, or 70%? I am trying to get a feel for it.

**Paul Lincoln:** Yes, salaries will still not compete with potentially double the salary available in the private sector.

**Baroness Doocey:** So it will be 50%.

**Paul Lincoln:** The kinds of salaries that we might pay, depending on the type of role, could be in the region of 50%, or less, than in the private sector. However, as the Minister said, the things that we can offer, such as the strong sense of mission, still attract people to come to work for us, as does the kind of data when it comes to data analytics, and the way people can utilise datasets. Our people can do things in that sense in a way they cannot do in the private sector.

**Baroness Doocey:** I understand that, and that in some ways it is a vocation, of course, and that is how we all want it to be. Could you also give me an example of retention? How many of your key people have been lost in the last, say, three years?

**Paul Lincoln:** I do not have any stats on retention, but we can write to the committee.

**Baroness Doocey:** That would be helpful. I am trying to ascertain how many people you are getting in. Are you getting the best and the brightest because they have a mission—they want to do it—but they then get to the stage of thinking that, despite wanting to do it and finding the work interesting, they cannot live on the salary?

**James Cartlidge:** There is a very important point to make about military personnel. In the Defence Command Paper refresh, the single biggest subject we touched on was the people who serve in our military. This is about many issues apart from pay, particularly when you are talking

about retention. I have responsibility for infrastructure and the estate. Among other things, it is about the estate, where people live, showing that we value them by investing in our accommodation. It has been documented that there have been significant issues with accommodation, but I am committed to improving it. It is about considering the structure of careers. There is the idea of zigzag careers. As we are getting into military careers, I might bring General Copinger-Symes at this point.

**Lieutenant General Tom Copinger-Symes:** Thank you for the question, because, frankly, it is a question of how we move faster here. I think your focus is mainly on specialists. I will come to some of the generalist issues too, because they are increasingly important.

Regarding specialists, the Second PUS mentioned the Digital Skills for Defence programme, which was signed off just the other day by the Minister, but was announced in the Defence Command Paper refresh. It looks at a range of interventions from bursaries at school, potentially, and we have an exciting pilot about to start focusing on cyber careers in particular, but it is part of this wider piece and, of course, AI is increasingly important in cyber defence, amongst other things.

On apprenticeships, to your point about the brightest and the best, we have a lot of apprentices in our workforce, as I think you know. The Army is the largest apprenticeship organisation in Europe. I think the second is either the Royal Navy or the Royal Air Force. Forgive me; I cannot remember which, but we can write to you on that. Overall, Defence just knocks the importance of apprenticeships out of the park by a mile.

Many of the brightest and best now working in industry came through Defence as apprentices, as did many of the brightest and best working in intelligence agencies. Although we may well be recruiting for those with aptitude and talent rather than baked-in skills, that is a hugely important part of the ecosystem that the Minister spoke to. Digital Skills for Defence is focusing a lot on that, as well as on graduate-level and PhD-level skills, particularly for defence and science and technology. That is the specialist area.

Looking at retention in that area, we are never going to compete financially, but I want to stress the learning and development piece. We are an amazing learning and development organisation, and that is recognised. I do not think we speak about it quite enough, if I am honest, but the chance to upskill in defence and then go elsewhere is important.

Sticking to retention, the mission is the thing that retains people, because it is awesome. Speak to anybody who works in AI in a bank; with the greatest respect to them, they get paid more but they will not tell you that they find their job fascinating every day. Speak to somebody working in defence intelligence on this; they find their jobs fascinating every day. Retention relies on that, but many of these folk will not stay as I have for 30-odd years in Defence. They will probably leave after five or seven years, but I do not think that matters. If they join a defence prime, a digital prime, a small or medium-sized enterprise or the

Intelligence Agencies, or go out into wider industry, that is still value-add for the nation, and I think that is part of our role in defence. I will stop there on specialists.

**Baroness Doocey:** That is very helpful. It will help us, particularly me, to understand how it works when I see all the stats about how long people stay, for example—you just mentioned five to seven years. It would be very interesting to see that for the different grades. That would be great. Thank you very much.

Q169 **Lord Hamilton of Epsom:** I was in the Ministry of Defence when the Levene reforms were brought in. As you know, Levene made the point that we should set up a requirement and keep to it. I have to say that it did not happen in my day, and it does not seem to have happened since, mainly because the military interferes with the procurement process as it is going along. A friend of mine on the Defence Committee in the Commons says that the Army is the worst perpetrator when it comes to this.

You say that you will keep a dialogue going with SMEs producing new ideas. I can see the advantages of that. I can also see extra costs and delays as a result of adding on to what you already have.

**James Cartlidge:** As I think I said in my opening remarks, I would have been very surprised if there were not some similarity between the issues facing a Minister for Defence Procurement some years ago and the Minister now. One of the challenges is complexity. Tacking on extra requirements and the search for a perfect platform and so on are well-worn points in this discourse. In most procurements, that is not true, or it is not true to such a significant degree that it causes the sorts of problems that are well documented. There may be very good reasons for it. Urgent requirements may arise, as came up in respect of Iraq and Afghanistan, but, equally, new information may come about.

There are any number of reasons. We are making progress, as shown by the figures from 2020-22 on the reduction in time, and, as Paul Lincoln said, in the plans we have set out in the Defence Command Paper refresh on reducing procurement times, but they all point to the fact that overall we have to be more nimble. Whether you like it or not, we cannot go on as before because our adversaries will be moving at pace and because the world is moving at speed. We have no choice but to improve things. The key will be the idea of trying to go for the non-perfect—I do not like using that phrase—the 70%, but then committing to spiral development once something is in service, which is a very effective way of bringing capabilities forward.

To your point about SMEs, I am not sure that it raises costs. That is more about general engagement. I do think that we need to engage with SMEs. I am very conscious of that. Just walking around one part of our very large and brilliantly attended stand at the Polish show on Tuesday were businesses 90% of whose income was from exports. Others' was a smaller amount and they were trying to bridge the gap. You are talking



about a huge variation. I, and we as a department, have to learn lessons to enable those companies to grow, to export more, and to find our procurement easier to deal with. There are good reasons, as you will know from your experience, Lord Hamilton, for why procurement in the MoD will always have some complexity. There are issues of secrecy, and so on. Ultimately, however, our businesses, and our industrial capacity, is the key to how we bring forward at pace the new platforms and technologies that we need if we are to compete with our adversaries.

**Q170 Lord Mitchell:** Good morning. I want to ask about training data. How does the Ministry of Defence ensure that it has access to unbiased representative training data for AI systems? Coupled with that, are its procurement processes capable of supporting the acquisition of software and datasets as well as hardware?

One of the things that come out in all the studies we have been doing is that suddenly, in the area of defence, where most of it has traditionally to do with hardware, we now find software very much to the fore. I am an IT person myself. Software seems to be all. How does procurement change its emphasis in being able to deal with software.

**James Cartlidge:** There are a number of points there. You are absolutely right on the general point that we are moving into a digital age. It has been happening for quite a while. There will always be an inherent focus on the big new shiny platforms and on tangible things that we can see out in the field or on the sea that give us photo opportunities, as I am sure other Defence Ministers have found, but you are absolutely right about the growing importance of things like electronic warfare, cyberspace, and space itself. This is the future where we need to compete with our adversaries.

You asked about synthetic training data, which I will bring Paul Lincoln in on in a moment. We have used synthetic data in a number of limited-use cases for the training of AI models and consider that it has a place in that training and can be used in combination with real data to provide a richer dataset for training rather than real data alone. However, there is some complexity in this point, and I wonder if Paul could try to put it into layman's terms.

**Paul Lincoln:** I will do that briefly, but I might ask General Copinger-Symes to give some specific examples.

As you have said, there is a risk of bias, which we have recognised in our ethical principles, so we need to be able to demonstrate not only amongst ourselves but more publicly the assurance that we are taking the potential for bias seriously.

The core thing, as the Minister said, and as you said, is that there may be occasions when we could be in a position where, because the training sets that were available on a military, specific set of scenarios may be lower than that which you see in, for example, a commercial, public set of scenarios. We might need to train particular parts of AI to use synthetic data to augment what we might already have.

**Lieutenant General Tom Copinger-Symes:** Just to expand on that, inevitably if I were training a system to recognise what a cat looked like, I would have the whole of the internet to trawl for data and could get a very broad and diverse dataset. If we are training a system to recognise a threat tank across the whole world, our existing dataset to train on that might be slanted to where we have operated previously or where our intelligence gathering has been focused. For instance, you might be looking for a tank but all the tank images are against a European, dark-green background rather than a desert, jungle or arctic background. To prevent that bias—we have a slightly different skew on bias from the bias we talk about in civilian life—in order that our system can recognise the tank globally we might have to create synthetic data, to create backgrounds that simulate the environment we will be searching in. That means that the whole system will be far more effective at finding the enemy tank wherever it is in the world. That is the point on training data.

Would you like to come on to what we are doing about acquiring data and software specifically, Paul, or would you like to come back in first?

**Paul Lincoln:** I would like to make one further point on that. In support of what Tom Copinger-Symes just said, we could also be looking for wider commercial data on civilian uses of military transportation so that we are quite clear that any form of artificial intelligence is making the appropriate distinction between the two.

Q171 **Lord Mitchell:** One final question. I did use not the phrase “synthetic data”, although I might have implied it. Could you define synthetic data for me, please? I think I know what it means.

**James Cartlidge:** I will attempt a layman's definition. The SME I ran was web-based, and I hope that gives me some understanding. I think the general explained this very well. Essentially, for a dataset not to be at risk of bias, which I think was the word he used—it seems a little pejorative, but that is what we are talking about—you insert into the dataset artificial creations that enable the system to perform more effectively. Does that make sense?

**Lord Mitchell:** Yes, that makes sense. Thank you.

**The Chair:** Would we be right to take from what the General was saying that if you are going to enrich your synthetic data with real data, that will make considerable additional demands of your intelligence resources?

**James Cartlidge:** Ironically, of course, that is precisely the type of application where AI may, in a very positive way, be very helpful. It is interesting that when you focus on AI, in the context that we are discussing, as soon as the word “weapons” appears, there are all kinds of connotations, and it is very important that we debate and discuss the various scenarios that could emerge, some of which could be quite serious. Nevertheless, it is worth stressing that AI will offer many opportunities and benefits, many of which will be relatively mundane—in the sense that they are not high-profile weapons-related issues—but

there will be many applications in head office, in administration, in routine deployment where AI will have a massive beneficial effect. One of those, of course, will be data management—management of the enormous quantity of data that we have. That is one of the most potentially positive uses of AI, and it is already real in the commercial world and, to a degree, in the military world, too.

**The Chair:** I am thinking also of the practice, which we have seen quite a bit of in the conflict in Ukraine, of disguising targets so that they become more difficult to acquire by any system, whether AI-enabled or not.

**Lieutenant General Tom Copinger-Symes:** That is exactly why it is so important to bulk out our real datasets with synthetic datasets.

Forgive me, but the recent upsurge in generative AI using large language models is a perfect example. To the Minister's point, even three years ago, to create large synthetic datasets would be mandraucic; a human being would have had to assemble those datasets. I have not done this, so I am taking a bit of a risk in saying this, but I am pretty confident that with generative AI now we can ask the machine to generate. We can give it a picture of a threat tank, give it a whole range of geographies around the world with a whole range of camouflages, and ask the AI to produce that synthetic data for us. It can do that very quickly and provide a much bigger dataset which we can then train our recognition models on and make them much more effective and prevent spoofing by camouflage and other deception techniques.

**Paul Lincoln:** This is not a UK-unique issue. This is where we can share datasets with allies and partners, so we are not just going off a small set. We can expand the real data that we have with others as well.

**The Chair:** I think you have taken us very elegantly into an area that Lord Triesman would like to explore, if he can hear me.

Q172 **Lord Triesman:** Thank you, Lord Chair. I can hear you very well, and I hope I am audible. Good morning, Minister and colleagues. Can I ask what systems the MoD currently uses in targeting and whether you benchmark those systems against non-AI-powered and human-operated systems. That is my first question. Secondly, is the Ministry of Defence working on any counter-AWS systems? I do not want to intrude into things that may be secret, but I would be grateful for a broad answer. To the Lieutenant General, may I say that there is not a day in investment banking that is not exciting?

**The Chair:** Can I just check that our witnesses heard those questions?

**James Cartlidge:** I did not entirely hear the first question.

**The Chair:** It was about what systems the MoD currently uses that use AI in targeting. The second—the complement, as it were—was whether any of those systems' performance is benchmarked against non-AI and human-operated systems.

**Lieutenant General Tom Copinger-Symes:** Thank you for the question and thank you for pointing out that banking is indeed a very exciting career for anybody who wants to go into banking.

The best answer to your question is that we have a number of systems in development at this stage, particularly for recognition tasks, but they are still in the upper ends of research and development rather than in full-scale deployment at the moment. Having said that, on the recent Sudanese NEO—non-combatant evacuation—we worked with industry to use some commercially available tools that were not finding the enemy but were finding safe spaces to move people to, finding gatherings of civilians and protected personnel and so on. So we are already using some commercially available tools, but the specific deep military application stuff is mainly still in the R&D area. Of course, as part of our value for money assessment as those programmes go on, as well as the safety and responsibility aspects we are comparing them with non-automated or autonomous systems. Inevitably, you want to check that this is delivering a bang for a buck compared with the old-fashioned system of having 10 imagery analysts looking at the same thing.

I hope that attends to what we currently have in service and how we benchmark them against current systems.

**James Cartlidge:** By the way, just to be clear, the UK Government are obviously deeply committed to having a successful financial services sector. It is one of our key industries. We have to have a way of paying for our military procurement, after all.

Further to my previous point, although your question was about weaponry, there are projects at development stage that are very significant for MoD that are in the military space but are not necessarily connected with offensive or defensive capability, such as Spotter, which is to do with image recognition and is a live project. We also have a project to do with our wheelbarrow. I do not know if the two former Ministers are aware of that term, but it is the term we use for the robot that goes out to areas where there have been chemical or radiological weapons and so on. By the use of a current AI project we are seeking to significantly improve its performance, which, by the way, underlines that this is not just about military capability in terms of war fighting; it is about tasks that free up humans from the mundane and, in particular, keep them out of harm's way, which is another potentially very positive use of AI. As I say, we are bringing those projects forward into development.

Q173 **Lord Hamilton of Epsom:** There was a report in the papers some months ago that the Americans were trialling an AI system and it basically went completely AWOL. It blew up the operator, killed him, and then blew itself up. I was interested to know what the Americans have done subsequently. They have denied that it ever happened. Let us hypothesise that this happened in the MoD. What would you do if that happened when you were trialling something on a simulator?

**James Cartlidge:** It is a very general point. You will appreciate that we cannot speak for the Americans and that we would not want to be drawn into hypotheticals on something as sensitive as that, as you will know.

**Paul Lincoln:** We are treating the way we adopt artificial intelligence into the department as a whole and into weapon systems the way we would treat any other type of technology. Although this is a novel technology, our approach is not novel. We have a series of tried and tested procedures for taking in novel technologies. We go through weapons and safety reviews, make sure that the technology is compliant with doctrine, and work through what the rules of engagement might be. We apply exactly the same set of processes from development and concept all the way through to service. A critical thing is that a safety case would have to be gone through as well as trials and tests as part of that, and when it comes to artificial intelligence, as with any other weapon system we would conduct an Article 36 review to make sure that it was compliant with IHR before we entered it into service.

**Lord Hamilton of Epsom:** You would send it back to the manufacturer in that case.

**Paul Lincoln:** Depending on the particular circumstances of a case, we would take whatever action would be appropriate.

**Lieutenant General Tom Copinger-Symes:** I would just point out that although we are talking about AI here, as the Second PUS said, sadly, that sort of thing happens with conventional systems, albeit rarely, but we have very well-practised procedures when something happens, whether it is a serious incident review or a service inquiry, and clearly we would stop using the fleet of similar systems until we had bottomed-out what the issue was. Although there is a certain excitement about a particular system, I do not know about the incident you are referring to, but we have tried and tested procedures for dealing with such incidents and working out the lessons we need to learn and how we take them forward to make sure that the system is safe and responsible again, not least because if you do not do that, your soldiers, sailors and aviators will not trust that bit of kit and will not use it, which is one of the reasons why we take such things so seriously. Above all, they will not sleep at night if they do not know that that bit of kit is achieving what we do safely and responsibly.

Q174 **Lord Browne of Ladyton:** I want to go back to training data and, in particular, General, to your answer. I intend to ask you what I think is a perfectly straightforward question but which of course in this area will turn out not to be. I will give you an opportunity to prevent people from inferring from your evidence something that you did not intend. Synthetic data is not a panacea for bias from data that has been gathered from the internet. It cannot be simply because it has been created by people who will bring in biases, and that is potentially exaggerated with generative AI, which itself may have been trained on data that had biases. So it is almost impossible to eliminate biases from training data. Is that not correct?

**Lieutenant General Tom Copinger-Symes:** I think you are absolutely right. Eliminating bias absolutely is impossible. What we can do is control for conscious bias, the biases we know about, and again I want to remove this from the way bias is spoken when talking about how AI is being used in society. This is about biasing a system towards recognising a particular tank in a particular environment rather than making it much more applicable. But yes, you can only control for that risk. Of course, what we hear about today is how we control for a whole bunch of risks in our use of AI, so complete elimination? No, I do not think you could.

Q175 **Lord Clement-Jones:** I was a bit surprised to hear the statement just made by the panel that you treated AI in the same way as other technologies. Without wanting to anticipate overly some of the later questions, you have made very great play of the ethics involved in AI acquisition and deployment. I do not see how those two statements, those two factors, stack up. Is there not a particular kind of ethical impact assessment when AI is deployed for weaponry? Is there no extra factor involved in AI acquisition and deployment?

**James Cartlidge:** The point that I think the Second PUS was making, which is a very good one, is that it is not new for the Ministry of Defence to have a period where it is looking at novel technology coming forward in the weaponry domain. That was the first point. Again, I also am not trying to pre-empt later questions, but inevitably we can go in a circle.

**Lord Clement-Jones:** It is a natural follow-on.

**James Cartlidge:** If you talk about the legal side, international law and so on, that is one of the reasons why we support the continuing use of international humanitarian law as the primary regulatory vehicle for AI and autonomous weapons, lethal autonomous weapons, fully autonomous weapons and so on, because of the consistency of application, which, by the way, therefore means that it is well understood internationally and there is widespread support for it.

I do not think we are denying that there are obvious specific factors to AI, but there is this important word, ubiquitous. It is not necessarily one particular thing. It will touch everything. That is why this point is so important. Be wary of being drawn into trying to be too tight in definitions, or too specific in policy. We are very confident about the architecture of international law and the way our ministry is subject to its own regulation, to UK law and so on. These are very important anchorings that have stood the test of time. You are absolutely right, however, that in technology terms AI will eventually bring many new challenges.

**Lord Clement-Jones:** So you are not denying that even currently—we are not just talking about the future—there are specific issues that may arise in relation to AI applications, particularly in weaponry.

**James Cartlidge:** Of course.

**Lord Clement-Jones:** That did not seem to be the case when the

Second Permanent Secretary was saying that.

**James Cartlidge:** Of course. That is why we are here. That is why you are holding what I think is a very important investigation. We recognise as a Government that there is public anxiety about artificial intelligence. That is precisely why the Prime Minister will be holding an international summit in the autumn about AI safety.

It is worthing caveating that defence will not be included in that. Nevertheless, it is a very important statement of the Government's overall commitment to ensuring that there is public confidence in the way we explore AI. However, as I have said earlier, that there will be many positive uses of AI. There already are. There will be many applications that enrich our lives, and in the defence context it will release personnel from mundane tasks or from those that could put them in harm's way. But yes, ultimately, we would be incredibly naive not to look at AI in the context of weaponry and military capability in the field.

**Paul Lincoln:** May I unpack my previous comment? I think we recognise, as the Minister said, that this is a ubiquitous technology. It has enormous potential the world over. Artificial intelligence has been used for over 20 years by companies such as Google. It would be wrong to say that we did not think the Government should be in a position where we should benefit from that.

We have set out clearly in the published artificial intelligence strategy that there is a series of benefits that we think apply to defence. It is not just in weapons systems; it is beyond. We also recognise that there is a set of risks associated with AI, which is why we have very clearly set out what our ethical principles are against those risks. Notwithstanding that, we have ethical principles, and we say when it comes to things like bias that we need to demonstrate as part of how we develop through the life cycle that we are mitigating that bias in the same way we would for humans or anything else.

However, there are still some constants in this process. We will still, as we adopt novel technology, be making sure that it is safe and that we have the right safety case. We will make sure that its use is context-appropriate, as we would with any other weapons system. We think that the most appropriate way of doing that when it comes to thinking about the law is international humanitarian law, as the Minister said, which has the benefit, of course, of being technology agnostic.

**Lord Clement-Jones:** That is much clearer. Thank you.

**The Chair:** Thank you very much. Minister, forgive me if I am mistaken, but I do not think we had an answer to one of Lord Triesman's questions, which is whether the Ministry of Defence is working on any counter-AWS systems.

**Lieutenant General Tom Copinger-Symes:** The general point to make is that of course we are, and you do not need an AWS system to counter

an AWS system. We have an existing range of protective defensive systems that apply to countering AWS systems. Of course, it is a race, so we will have to keep pace with that race, but we have existing systems that would help us deal with any AWS systems that we found on the battlefield. We are investing so much money in development in order to keep pace with that competition.

**Q176 The Lord Bishop of Coventry:** Minister and colleagues, thank you very much for your presence and evidence. You are reinforcing a reassuring theme, which has been evident in our deliberations, that AI is a good servant and needs to be harnessed. It carries risks, so we do not want it to be the master, and that applies to AI in autonomous weapon systems.

In our sessions, we have heard on that theme a lot of fascinating and somewhat tantalising information about the human-machine interaction, human-machine theming, which of course relates critically to questions of context-appropriate human involvement, to IHL and to those ethical principles. Mr Lincoln, you told us a bit about accessing Google training. I would be very interested in anything more you could tell us about the training the MoD provides for the particular demands of working with autonomous weapons systems and whether that includes involvement in systems design, development, risk profiles, and—this is where I find it deeply tantalising—training for weaponry that is learning and developing in deployment. Then, of course, there are the particular psychological pressures. It is a huge subject, but can you say more about the training that is provided?

**James Cartlidge:** There is quite a bit to unpack there. I will come back to Paul Lincoln on training if that is all right.

To the point about context-appropriate human involvement, this is probably one of the most important phrases that we will discuss in the committee. I think it is helpful to use a real-life example. The Royal Navy uses a gun called Phalanx that contains in its potential use a capability that can arguably, for part of its use, be described as partly autonomous/automated. The crucial thing, however, is that it can operate only if there is the context that we are talking about, which is appropriate and that must involve human actors. In other words, to put it as simply as I can—I appreciate that I am now talking about military capability here—it has to be switched on because there is a situation that presents a threat or circumstances where the commanding officer believes that its use is right and appropriate. This is crucial, because it is that anchoring, that context-appropriate human involvement, that ensures that that system is compliant with international law.

I hope that gives you a tangible example. If that did not happen and the weapon could just choose to come on and engage, that would be a fully autonomous weapon system, but it is precisely because of that context-appropriate human involvement that it is not fully autonomous and is compliant with international law, and it is an important piece of capability for the Royal Navy.



**The Lord Bishop of Coventry:** Thank you very much, Minister. I think I understand that. The human-machine interaction in that case looks relatively straightforward. This is probably oversimplifying it, but I guess what I am thinking about when AI is added to the Phalanx and to other systems is that it seems to take on an exponential complexity. How do you meet the training needs of the operators?

**James Cartlidge:** I will bring in Paul on the training and then General Tom on the morale issue that you also raise.

**Paul Lincoln:** We might box and cox. The Minister mentioned context-appropriate human involvement, and that goes through the whole life cycle. It is not a specific part. We need to think about this from the design, the policy, the research and development, the risk and where the human involvement might be, whether it is in the policy-context setting, whether, as the Minister just described, there is an on/off component to it, or how we set the rules of engagement and so on. That is no single set of things. It is an iterative process that we will go through.

Ultimately, however, as you describe, there is a human/machine interface that comes into play. That has been true throughout history when it comes to any sort of weapons system, whether it is a rifle or something more complicated. This is about how we harness the way the human can make judgment calls about how they will prosecute some form of action. In doing that—the Minister mentioned Spotter, for example, as a potential recognition tool—we try to make sure that we provide more reliable, more capable information and analysis to a decision-maker so that they are more capable in their jobs. You can imagine a situation where you might be cold, wet, tired and hungry. As a human, you are more likely to make mistakes in those circumstances. Having provided you with greater analysis than you might otherwise have had to support your decision-making is a positive in that situation. Artificial intelligence in those sorts of circumstances will, of course, be very helpful.

I go back to the training, and General Tom will no doubt want to come in on this. How do we make sure that people are trained in their weapon systems? That, again, is a process, another set of things that we will do. People need to understand that they will be trained and will be confident in the use of the systems they have, and we as an organisation need to be able to certify that they are safe and reliable in the way they are operating.

Although this is not the position we are in, if there were a situation where you had machine learning deciding to do something substantially different, we would have to go back through the Article 36 process, which would say, "Is this still compliant with the original purposes of that weapon system?" We cannot enter weapons into service unless they have gone through that process.

**Lieutenant General Tom Copinger-Symes:** As with one of the previous questions, I want to tread carefully between saying, "It's all right. We've seen all this before, and we've got processes" and

recognising that with AI we constantly need to check whether our processes are good enough. Let me start by saying that we have seen this before and we have processes. Whether it is training individuals or training them with their kit and then building up from small groups to larger groups, that is what we do every day of the week in Defence: train our individuals, our teams, with their kit and build up increasing competence and readiness so that we can respond. To that extent, this is normal business.

Of course, we need to be watching where AI changes that. The really fascinating bit here is that, whether it is a recognition system like Spotter or, and no doubt we will get into this, a weapons system, AI is itself now able to learn—so machine learning. If we do see something that is outwith what we can currently do, that is where the machine is learning. But, as the Second PUS has said, we already have safeguards, such that if the machine were teaching itself to do something else, we would have to run it back through exactly the same cycle that we have now to make sure not only that it is safe—in other words, it does what it means to do and it is not going to harm the sailor operating it—but that it is responsible and obeying our ethical principles.

That is the operating bit. I do not want to sound complacent about it and say, “Oh, it’s all right. We’ve always done this”, but we do—

**The Lord Bishop of Coventry:** I have another ignorant sort of question, forgive me, but how long does an Article 36 review take in that sort of situation? Is it quick? Is it long?

**Lieutenant General Tom Copinger-Symes:** That is an excellent question that I think we will have to come back to you on.

**Paul Lincoln:** We can come back on it, but it is case by case, depending on what the system happens to be.

**Lieutenant General Tom Copinger-Symes:** Clearly, in the past, for conventional weapons systems that has tended to be for something that happened before deployment, and we are already looking at in-field assessments—I think that is the technical phrase. Where a system is learning in the field, we are already looking at how we would adapt our current processes to be able to just check on that in the field, not just back. We are probably not there yet, but that is just to reassure you that we are absolutely thinking about it.

Commanders were also mentioned. Operators are clearly important, but commanders are hugely important to this, because most of us did not grow up in a world of AI. We have a range of interventions there to make sure that commanders think about this. Once every six weeks or so—the Second PUS will have been there on Monday—we have a regular programme whereby the most senior military and civilians in defence come together. We get technologists, commercial entities, our own scientists in to sit down with them and think, “What is different about how we lead, how we run Defence, in a digital age?” A lot of that, of

course, focuses on AI, although not exclusively AI, so we are understanding through the whole life cycle of Defence what this means for us. That is not to say that it is done and dusted, because everything is changing. Pitching to the CDS and the PUS at the time that our most senior leaders need to learn was quite a difficult thing to do. It was pretty impressive that they agreed to it, and, as I said, we get the most senior leaders coming together for a learning session every six weeks or so.

**Q177 Lord Fairfax of Cameron:** I want to ask about the stricture of context-appropriate human involvement. You mentioned the well-known example, of Phalanx, which is a defensive system, so my question is, I hope, a straightforward one. Is there perhaps therefore a distinction to be made between offensive AI systems and defensive AI systems so that a defensive AI system might be more tolerable or acceptable, for example, or considered as such, than an offensive one?

**James Cartlidge:** The legal position is absolutely clear. Whatever the capability, it must comply, we must comply, with international law. That stricture on context-appropriate human involvement must apply. The Phalanx is a good example because of the sense in which it is partly automated. You could even use the example of an anti-tank mine, which, once it has been laid, will be triggered by an event. You could argue that at that point it is autonomous. Or is it automatic? Whichever word you use—maybe that is semantic—the point is that there is no direct human interaction at that point. Paul rightly used the phrase “life cycle”. You can look at that across the whole life cycle. The landmine was designed for a very specific purpose, which means that it complies with those key long-standing principles of international law on whether, for example, it is discriminatory, proportionate and so on. That is the key consideration, and that is where the defensive or offensive comes in.

**Paul Lincoln:** I think there is a risk of drawing a binary conclusion here, rather than saying that there are four principles that we apply—proportionality, necessity, distinction and humanity: in other words, principles that prohibit unnecessary suffering. Those are the context-specific parts to that. The other thing that is context-specific about the example the Minister also gave about anti-tank weapons is that in a military situation we, as the UK, would put up signs saying that that is a minefield. We also have rules of engagement as to how you can deploy those systems, and so on. As long as we are compliant with international humanitarian law, that is the basis on which we would start the conversation.

**Q178 Lord Fairfax of Cameron:** If you had a system that was compliant with international humanitarian law but there was no human involvement, would the lack of human involvement trump compliance with international humanitarian law?

**James Cartlidge:** You are implying that it is fully autonomous, essentially.

**Lord Fairfax of Cameron:** Like Phalanx is once it is turned on.

**James Cartlidge:** The very fact that it has that context of human involvement means that it is not a fully autonomous weapon; it has elements of autonomy. This is an incredibly important distinction. When I was asked about this at the last session of Oral Questions by Theresa Villiers, my colleague, I was absolutely clear, as we have been in all the answers we have given on this, that the UK does not possess fully autonomous weapons systems and has no intention of doing so. That is the key difference. You could clearly have something that has elements of autonomy or automation, as we have described. There are other examples—Brimstone, for example—but the key legal point is whether you have that ultimate human involvement, because it is that that ensures that it is grounded in the principles that Paul was talking about.

**Paul Lincoln:** Of course, Phalanx has been programmed with specific sets of scenarios under which it will operate once it is turned on. How it will operate is designed by humans. Thinking about how the overall system works, the key thing about a fully autonomous weapons system is there would be no human accountability in those circumstances. The Government have been very clear that humans will always be accountable for the actions they take, however they may be enabled.

**The Chair:** We are likely to come back to IHL a little later.

Q179 **Lord Browne of Ladyton:** I think this question is for you, Mr Lincoln, because you raised the Article 36 review. I want to ask you a very straightforward question. Is it possible for a technology with a black box—in AI terms, that means a system where even the person who built the technology does not understand what is going on inside it—to pass an Article 36 review? Also, when will we return to publishing the outcome of Article 36 reviews so that people more broadly can have trust and confidence in our weapons systems, which they used to have when they were published, as the Americans do?

**Paul Lincoln:** On Article 36 reviews, as I said earlier, which we have to conduct on any new weapons system coming into place, there is a set series of steps. This is undertaken by the legal teams and wider teams in the Development, Concepts and Doctrine Centre down in Shrivenham, which leads on that work. It is internationally renowned for that work and indeed, as part of that, will teach courses in Geneva on this. If it is helpful to write to the committee with more detail, we can do so. As part of that, as we talk through the life cycle, people need to understand how these operate and the safe and repeatable steps that those systems will take in order to pass the use cases that those weapons systems take to be allowed to go through that end process and come into service. It has to be able to demonstrate that this is reliable.

**James Cartlidge:** You asked about publication. Essentially the position on publication is subject, as ever—as it would have been when you were Secretary of State—to the provisions of national security and so on.

A note from our officials about the question on time. I am not sure we have an actual database of average times for Article 36 assessments, for

example, but we will look into that. Essentially, I am assured that an urgent capability requirement could happen very quickly indeed, but a complex weapons system could take far longer. It varies, and it depends on the circumstances of the specific weapons system in question, but I will see if there is any hard data on timing.

**The Chair:** Perhaps when you write to us with the note that the Second Permanent Secretary promised on the steps and the detailed process, you could illustrate it with some life-cycle cases of assessment of weapons systems. That would be very helpful.

**James Carlidge:** We will endeavour to do so.

**The Chair:** Thank you very much.

Q180 **Lord Grocott:** Just getting back to the Bishop's original question, the shorthand journalistic summary of it would be "killing at a distance", I think. I would like your thoughts on where the person doing the killing—I am trying to avoid it being too much of a journalistic heading—is far removed from the battlefield. Maybe I can characterise it as someone getting up in the morning somewhere in Britain and going to work, metaphorically pressing a few buttons and taking out a building 3,000 or 4,000 miles away, maybe with people in it. I know we are talking about semi-autonomous systems, but this is a new type of responsibility and challenge, I would have thought, for the people doing that kind of work. Are there any special training facilities that address those particular issues? Was it the Minister or Mr Lincoln—I cannot remember—who said that the soldier may be cold, wet, tired and hungry when he is making a decision, but here is someone who is warm, comfortable, safe and so on? How do you approach whatever mechanisms there are for assisting, guiding and training people in that situation?

**Lieutenant General Tom Copinger-Symes:** Thank you, because it is a very important question at the moment, of course. It has nothing to do with autonomy per se; it is all about remoteness. Indeed, in the last few years, this has come into stark relief. Generally, in terms of pastoral care, for instance, we would scale padres and chaplaincy, in the Army anyway—towards deployable folk. The thing that triggers the need for the chaplain is those who go overseas to fight. As you are suggesting, there are people who are involved with that fighting back home as well. This is not just about kinetic operations; it is also about cyber and other sorts of engagement, where people can be under very significant stresses but going home to their family at night, rather than being wrapped with all the stuff we have developed for hundreds of years to keep people feeling part of a team and connected to the mission.

I have given you clues for some of the things we think about. How do we make sure that pastoral care is extended right back to where people are engaged in what we used to call the home base? That is a very important part of it. How do we make sure that all the ways in which we use teamwork to wrap people with care and comradeship and all those things extend to what we used to call the home base now? Those are the typical

soldier/sailor/aviator approaches to how we manage stress over the years, which we have optimised in the past for doing it 'over there'. Now we need to work out how we bring it back here. Some of that is about training, some of that is just about pastoral care and companionship, and some of that is also about leadership.

Of course, the whole history of warfare is that it used to be up front and personal, including the leader, the General or whatever, and now that is becoming increasingly remote. We have learned a lot over the past few years, but some of that is coming into even starker relief now and we are responding to it in the ways I have laid out.

**Lord Grocott:** I am sure we can learn from our colleagues in other countries and within NATO et cetera. Have any lessons been learned as to how to deal with this new kind of competence? Like everything else, as you have explained, nothing is new. There are elements of what has gone on before, but have we learned from allies?

**Lieutenant General Tom Copinger-Symes:** One of our advantages is that we have lots of friends. Most of our adversaries do not have many friends. The Second Permanent Secretary made the point that we can share data and therefore make our systems more secure. We can also share lessons. Whether it is in a NATO framework or a Five Eyes framework, we do share those lessons. I cannot point to any specific thing, but I can happily come back to you about how we are sharing lessons. So, yes, that is a feature of what we can do and how we can learn from each other. It is one of the great advantages of being an alliance and having lots of friends, as opposed to our adversaries, who tend to have fewer friends.

**James Cartlidge:** Can I give what I think is a very powerful example of that? Quite early on after taking this job, I flew to Benbecula and attended NATO's exercise, Formidable Shield, off the Outer Hebrides. It was very instructive for me that there was a live firing of a cruise missile essentially from the beach out into the sea, where there were naval ships from many nations, NATO allies. Crucially, as that missile was deployed on the screen, we could see the data of its journey and all the data about its potential impact being shared between naval ships from different nations: Italy, France, Spain, the US, Scandinavian countries and ourselves. As to General Tom's point, it is very powerful. There may have been, dare I say, a fishing vessel somewhere in the vicinity with some funny aerals on it watching that. I think they would have watched it all and thought, "These guys have a lot of friends".

That is the power of the deterrent of NATO, of having a large alliance, but you are absolutely right that, for it to be effective, we have to be able to share data. I think that is one of the best examples of it, but we will need to do so on AI in a range of fields, and of course we do that through Five Eyes. AUKUS will be another. The key thing for this country is that we often talk about our defence as if we are acting alone. In almost every context we can imagine, particularly militarily, we will be acting with

other partners and sharing data in the way I have described, to some degree.

**The Chair:** I think Lord Browne wants to explore the operation of CCW.

Q181 **Lord Browne of Ladyton:** I have spoken twice and not thanked you, the General or Mr Lincoln for your evidence. This is very important to us, and thank you very much for being here. The helpful letter you sent us yesterday sets us up for the questions I want to ask. In the penultimate paragraph, you write, "The UK is also a leading voice in international dialogues around the safe use of AI and autonomy in the defence domain. We are championing our approach to responsible and safe AI in forums such as the UN Group of Governmental Experts on Lethal Autonomous Weapons"—which, I add, is commissioned itself by the CCW to take this forward—"NATO and through other international partnerships"

Discussions on international regulation of autonomous weapons systems currently are stalled in the CCW, where countries have been considering this issue for decades and the only decisions they have ever made are, "We'll come back to it again in the next meeting". In recent years, Russia in particular has had the forums consensus rule and has used it as an effective veto on any meaningful activity at all, so much so that support for a legally binding instrument to regulate autonomy in weapons systems is growing. More than 90 countries have now expressed a position of favour on such action. There will be a resolution to that effect at the UN General Assembly this autumn, which is imminent.

When we were in our presidency of the Security Council, we had the first-ever debate on the risks and opportunities of AI, which was a very good thing. At that meeting, we gave a platform to the UN Secretary-General, which is a very good thing, who called for states to negotiate a legally binding instrument to regulate autonomy in weapons systems to be concluded by 2026, which I think is a very good thing. Now that that issue is being discussed in New York, with states discussing it in the UN General Assembly, my questions are: what are we planning for pursuing this issue in the General Assembly, if anything at all, and are we collaborating with states to raise the issue in the GA's first committee this autumn? Will we be building on our Security Council focus that we created on responsible AI by including the issue of autonomous weapons as a significant strand of the planned summit on AI safety this autumn? I promise you I will go into supplementaries.

**James Cartlidge:** Thank you very much. I respect your expertise in this. I know that after you were Secretary of State you continued to do a lot of work in respect of non-proliferation and so on. You will be deeply familiar with these international arenas in which we engage. I think the UK is a leader. All this work is multilateral, by definition, and we are deeply committed to that work. We have found the CCW an effective forum in which to pursue engagement on these subjects, because, above all, when it comes to international law, you have to try to build consensus. Russia, inevitably—especially in the current circumstances—is a country that will

present issues and challenges, but I think a lot of progress has been achieved.

You will know that, in October 2022, 70 nations, plus us, signed the joint statement on LAWS—lethal autonomous weapons systems. I understand that this year the LAWS sessions at the UN have been progressive and progress is being made, but the specific context now with Russia is a challenging one. We have many different fora—I think that is the correct way to say it—in which we have this kind of multilateral engagement, from the Five Eyes to UNESCO and NATO. NATO is obviously critical in the current context, so I am confident that we are doing everything possible to engage. We provide experts to the GGE, and I can assure the committee that we will continue to engage multilaterally, as far as we can, on this precise subject.

**Q182 Lord Houghton of Richmond:** Thank you, Minister, for being here. I put on record a thank you to the Second Permanent Secretary for helping to facilitate our forthcoming trip to the PJHQ, where I hope the committee will learn a bit more about the regulatory framework on the battlefield and the operational use of AI systems.

To an extent, my sandwiches are eaten and the fox is shot, because this is another question on regulation. Given the fact that by the time we conclude our business in November we hope to have made some recommendations, one of the things that might inform those recommendations is a recommendation for the optimum regulatory framework for the whole business. Perhaps a first strategic question—you have almost answered it—is the degree to which you have an established concept for the holistic regulatory framework of AI. From what you are saying, it appears to me to be rather no change. There is international agreement on certain conventions, there is national-level regulation, there is Article 36 and things like that, which do or do not allow things into service, and then there is a level of operational regulation, which is the responsibility of the battlefield and the level otherwise of meaningful human integration with that.

The majority of the evidence we have taken along the way presents hurdles for exploitation. Some of them are legal, some of them are ethical. Some of them, though, are in the area of technical uncertainty, about the outcome when AI is fully exploited. I wonder, therefore, whether there is anything you would wish to offer us that we can recommend back to enhance the regulatory system. It strikes me that there will be some iterations around Article 36 and what is or is not dangerous to bring into operational use, particularly with systems that—who knows?—may start to learn things only once fully deployed if we are to truly harness the technology. Some of the technological evidence we have heard is from people who say that it will be impossible to simulate the outcomes of some things that this AI will be able to do.

The soldier in me wants to think or hopes that the level of the regulatory framework on the battlefield will not be too constraining and that we fail to maximise the benefit of these systems. Indeed, it must be the case



that the maximum benefit to be derived from these systems is when you are using them at the maximum risk of what this technology offers you. I will run the fear that we might run scared of it. Dare I say, so much of the erudition in the Defence Command Paper is about the sort of alchemy of technology and yet we will fall foul of being scared of it. In a world where our opponents, many of them non-state, will not give a toss about that, will not give a damn about that—

**The Chair:** It is all right, I think it is still parliamentary language.

**Lord Houghton of Richmond:** —it will become self-defeating. I know that for all sorts of reasons the department has to have a public line that the ethics, the law and the degree of technical risk must be well-bounded, but does this not deny us the very advantages to an extent that we want? Are there not things that you should be recommending to us to cut some sort of slack in this regulatory framework that permits us or does not deny us the remarkable opportunity that it offers for military advantage?

**James Cartlidge:** It is an excellent question and, to be absolutely clear, as far as I am concerned we must in no way act naively or put restraint on our country's ability to exploit AI within the bounds and parameters of international law, but must act in a way that ensures we stay ahead of our adversaries. We should be in no doubt about that. We only have to look at what is happening in Ukraine. That there is some intelligence, potentially, about AI use by Russia, for example, in our defence artificial intelligence strategy paper. Irrespective of that, in a situation like this where you know that it is operating in a fundamentally nefarious way—very explicitly, it has invaded a sovereign country—there has to be a strong presumption that it will be pursuing investment in R&D and technology. Other potential adversaries will, and, as I say, there is intelligence that to some degree confirms that. We must not restrict our ability to respond, but, equally, we must operate within international law. It is a balance to be struck.

If you are talking about potential recommendations, it is not the role of a Minister to come to a committee and suggest potential recommendations, but I would hope to appeal to you to recognise that in government we have to strike that balance. If you look back in history, as Paul said, in some respects it is not a new situation for us to be faced with novel weapons. Throughout the time during which we developed the nuclear deterrent, we as a country complied with international law. We are a strong member of NATO. We have fought in campaigns to liberate countries and, as far as we are able, to do the right thing. Nevertheless, we have maintained our investment in technologies that give us the ultimate deterrent against the most extreme threats. That principle must apply to how we go forward with all technology, including artificial intelligence.

**Paul Lincoln:** There are two points from my perspective on that, one of which goes back to Lord Houghton's point about whether you can do things in the context of a conflict where we are trying to maintain our

strategic advantage. I think the benefit of international humanitarian law, if we go back to that, is that it balances the different components of necessity, proportionality, humanity and so on, as we do that.

The second point touches on one of the other questions about reliability. One of the things that we have set out in our published ethical policy statement, which accompanies the AI strategy, is the fifth principle, which is about reliability. I will just quote, if I may, from that. It talks about being “demonstrably reliable, robust and secure”. This is about meeting a purpose: “The MoD’s AI-enabled systems must be suitably reliable; they must fulfil their intended design and deployment criteria”—which, of course, can be about strategic advantage—“and perform as expected, within acceptable performance parameters”. Then we go on to talk about how we assure and assess that.

**James Cartlidge:** You spoke about something that is incredibly important, which is essentially the training of soldiers in these matters. I want to reassure the committee that is very important, but I think it is best if General Tom explains what happens on that front.

**Lieutenant General Tom Copinger-Symes:** You will all know that at the very lowest level, from the very first day of training, we bake ethical and legal training into a soldier’s experience. The one thing you cannot do as a junior officer, which I was not so long ago, is delegate the one lesson that you have to teach on the law of armed conflict and ethics. I hope that demonstrates just how baked in this is to the very fundament of what we are. Frankly, for those of us who have been around for a while, we have all made mistakes, but the reason why we sleep at night is that we did everything we could to make sure that those mistakes were minimised.

To build on that, international humanitarian law does not ask us to lose. It says that winning has to be part of that legal framework. We are determined that we will give our soldiers, sailors and aviators the very best tools they can have to win. The very spirit of experimentation and learning lessons, as Lord Houghton knows, does not stop in the training field; it goes on to operations and repeatedly goes through operations. So every mission is followed by a mission debrief, where we learn lessons and see where we could do it better and where we will change the next time we go. That spirit is alive and well, but it needs to be even more alive and well.

We need that spirit of experimentation to energise the home base and our procurement in the home base. That spirit of failing fast—in the words of agile folk—learning lessons and then going again and improving typifies what we do in the battle space. We also need to bring that home and make sure that we are moving at that sort of speed with that approach to risk that we have when we are our best.

**The Chair:** I think you are right, Minister, in telling us that we should not expect Ministers to volunteer our recommendations, but Select Committees do find it very helpful from time to time to be told, however

subtly, which doors might be ajar.

**James Cartlidge:** You have great experience in these matters in many other ways throughout your career. I was simply trying to make what I think is one of the most important points, particularly from the view of the Ministry of Defence, which is that we are proud of our record as a country. We have spoken about engagement in the UN, complying with international law and being a champion of international humanitarian law, but we have an ethical duty to our people and to those who serve in our Armed Forces to defend our country and to ensure that we maximise our pursuit of those capabilities that enable us to outcompete our adversaries.

**The Chair:** Thank you very much. I will simply leave that thought with you.

Q183 **Lord Fairfax of Cameron:** This is just a quick question. I was thinking of when might be an appropriate time to ask it, but maybe it is now. The Americans made a big announcement last month about their new Replicator system. I do not know whether any of you are familiar with that. This is basically about drone swarms. They made an announcement they could not compete with China in mass weaponry, so they might have to do it on a micro level, particularly in relation to drone swarms. Because of Ukraine, one gets a feeling that that might be the way a lot of this stuff may be going. Have you had seen the announcement about Replicator and whether we will subsequently find ourselves going down that route with individual mini-drones, which might be individually controlled or controlled in a cloud? I am just looking for any thoughts.

**James Cartlidge:** That is an excellent question. I did see it; I did not see the detail. I receive a large quantity of submissions every day, which occupies my time in terms of UK policy.

**Lord Fairfax of Cameron:** I was not trying to trip you up. I just wondered if you had seen it.

**James Cartlidge:** No. Again, I think you make another crucial point, which I have stressed repeatedly, which is that we just happen to be at a point in history when we could have been faced with a country that has the world's largest population or thereabouts, with potentially its largest economy or thereabouts, and extraordinary growth in its military prowess, but luckily it is just at the point where uncrewed systems will matter far more, where technology will matter far more. Therefore, going back to my previous point, we need to be invested in these technologies, because that it is our way of compensating, to a certain extent, for not having that degree of mass, because mass can be achieved now through uncrewed systems, through technology. I think that is what the US was stressing in its public comments on that system, notwithstanding that I do not have the precise detail of it.

In terms of drones and uncrewed systems, I think you are right. This is an area where AI could play a very important role. One only has to look

at what is happening in Ukraine and the extraordinary way in which uncrewed systems and drones have been deployed. Without wishing to go into too much about it, we have seen significant evolution, as you would imagine, in a live battle. We have seen significant evolution already in the way they are deployed and used, and we have to learn from that. You are underlining to me why this balanced approach of complying with international humanitarian law and reassuring the public about the safety of AI is incredibly important. Particularly for the Ministry of Defence, we nevertheless have to exploit these technologies as far as we are able to within the bounds of international law so that we can maintain our competitive ability with our potential adversaries. Drones will be a key part of that, I am sure.

**The Chair:** Can I just apologise for the absence of Lord Browne and Lord Houghton? We are experiencing competition from a debate on the Armed Forces in the Chamber, so that is the explanation for that.

**Lord Clement-Jones:** We will come to ethics later, but I just wanted to clarify something that you said, Minister. You said that it was important to stay within the bounds of international humanitarian law. By that, I assume you mean that that is where the Article 36 review comes in, and so on. I want to clarify, because you have stated that there are AI ethical principles involved here. I would say that it is not just about safety but about many other aspects, but there is a second process of evaluation for the safety of the public and so on. I just want to make sure, because earlier we had the slight confusion that this was just any old technology, but this set of ethical principles, which we will talk about in a minute, is specific to AI.

**James Cartlidge:** Of course, important in relation to that is that we have our own AI ethics panel. I am pleased to say that sitting to my left is the chairman of that panel, which is very handy. Can I turn to him?

**Lord Clement-Jones:** We will come on to it.

**James Cartlidge:** Okay, fine. Noted.

Q184 **Lord Hamilton of Epsom:** My question predated the one we just had and is about defining AWS. You indicated earlier, Minister, that you find it very difficult to define AWS. Does that matter? Can you regulate AWS without defining it, both nationally and internationally? Is the existing regulation adequate to cover AWS anyway?

**James Cartlidge:** That is a very important question. For clarity, it is not necessarily that it is difficult to define; it is just that if you try to seek a very rigid definition, that could have unintended consequences. That is how I would put it.

I want to be clear about the difference between autonomous and fully autonomous. Earlier this year, we sent a joint submission with the US to the UN about how we could take forward engagement on these matters. In that document, we referred to the US definition of AWS. Its definition, if it is okay for me to read it to you, is: "A weapon system that, once

activated, can select and engage targets without further intervention by a human operator". It includes, as a subset, "human-supervised autonomous weapon systems that are designed to allow human operators to override operation of the weapon system, but can select and engage targets without further human input after activation". In other words, outcomes are determined by the configuration of the system as a whole, not by the maximum autonomous behaviour that a system might be capable of.

Apologies for that being quite wordy, but the point I have tried to stress when I have been on the record—the main time I did this was in the last defence Oral Questions in the House of Commons—is that, for us, it is about fully autonomous weapons systems. This is really where the public concern is, because we can talk about these technical definitions. Dare I say it, I think the clichéd image that comes to the public's mind and causes anxiety is the idea of some kind of robot with complete freedom, heavily armed, sort of rampaging with no control and threatening us, which is, we hope, a long way from being a realistic prospect. It is why this is about the issue of full autonomy in particular: because of the issues we discussed relating to context-appropriate human involvement. That is where we are explicit in saying that we do not possess fully autonomous weapons systems and have no intention of doing so. I hope that helps.

**Lord Hamilton of Epsom:** Do we need extra regulation?

**James Cartlidge:** Our position, as I and my colleagues have said, is that we strongly believe that the current position with regard to international humanitarian law is the best way to regulate this internationally. There are many reasons for that, but perhaps one of the most important is the degree to which it is understood and there is consensus behind it. The processes involved in the development of international law are consensual and multilateral, but something to stress is that those principles have been around for potentially more than 150 years, because, of course, they are principles of conduct of warfare rather than regulating specific weapons systems et cetera. They have a long-standing universal application, which would apply largely to an autonomous weapons system.

**Lord Grocott:** Are you saying that there is no such thing as an autonomous weapons system—that all of them have degrees of autonomy? Is that part of the reason why you do not feel it right to give a precise definition—that there is nothing that cannot be coped with based on this being, maybe rather dramatically, essentially a development of things that have happened in the past?

**James Cartlidge:** The key point to stress is that there is no internationally recognised explicit definition, as I understand it, of autonomous weapons systems. We are talking about international law and international consensus, so it is important to recognise that. As I said, we think that the US definition is helpful, although that is not about fully autonomous weapons systems; it is about autonomous weapons

systems. Without wishing to go too far down the rabbit hole of these potentially quite technical points about whether there are systems that are themselves autonomous, I think that would be for other experts to say. I have referred to examples that we have had, which are those that have had elements of autonomy but which clearly, in many ways, are anchored in human control through the life cycle from their design.

As Paul said, Phalanx, for example, will be programmed with very strict parameters about the type of target it would interact with, were it to be used et cetera. To what extent there may be systems that could be described like that I would leave to other experts, but there is no internationally agreed definition. I think we can agree that for most countries there would be profound concern about fully autonomous weapons systems. I think they would be regarded, by their very nature, as not compliant with international law, precisely because—this goes back to our original discussion—they do not have embedded in them that context-appropriate human involvement.

**Paul Lincoln:** It is probably worth building on what the Minister just said. He utilised the US definition there. That is in the context of a US policy, which is that there must be context-appropriate human involvement as part of the deployment of such weapons systems.

**Lord Clement-Jones:** I understand the difference between the practical question and the definitional question, but I thought that NATO did have a definition of autonomy. What had not been done was putting the two things, the weaponry and the autonomous nature, together in a definition, but I think NATO has adopted a definition of autonomous.

**James Cartlidge:** Pass. I would have to look at that. I think the crucial thing regarding international law is that the consensus is about sovereign nations agreeing. That will include many countries that are not NATO members. I do not know about NATO. I do not think there is.

**Paul Lincoln:** I do not know about that, but there is a forum, the Data and Artificial Intelligence Review Board, which we are a leading member of, which deals with these sorts of issues. I think this is partly about the purpose of that definition. From our perspective, in principle we think that having a definition is not necessarily a bad thing, as the Minister said, as long as it is not overly constraining, because it helps the conversation about a seriously important subject. I think part of the debate is to what purpose people want to put the definition.

Q185 **Lord Fairfax of Cameron:** I will ask the next question, which is about the use of autonomous weapons systems by non-state actors—bad non-state actors, probably. I just wondered if you had any comments on that. We are already defending against some of that sort of territory.

**James Cartlidge:** You can appreciate probably that I am always very inclined to refer to examples that are already in the public domain. There is another one in our defence artificial intelligence strategy, which we published in June last year, which refers to "Non-state actors ...

repurposing commercial technologies to enhance their capabilities". It gives an example from 2018: "Iranian-backed Houthi forces in the Yemen used an Uncrewed Surface Vehicle, a USV"—probably a converted fishing skiff—"loaded with explosives to severely damage the Saudi Arabian naval frigate, AL-MADINAH. They have since used USVs to attack maritime targets ... While current systems may be limited to GPS waypoint navigation to reach a predetermined target, irresponsible proliferation of AI technologies could enable proxies to field far more dangerous capabilities in future".

Again, to the point that there is some intelligence on this, but a bit like my earlier point on state actors—we have one in particular in mind, of course—you have to work on the assumption that this could get into their hands and some of them will be working on it. That is unfortunately quite a serious threat and something we look at across government, because of course other departments, such as the Home Office, have equity in that, and there is a lot of work happening looking at that. It is very much a co-ordinated effort across government and with international partners.

**The Chair:** In terms of non-state and rogue state actors, the balance is tipping quite a bit, is it not? If you are looking at weapons-grade plutonium, you can see the plant to produce that from space. If you are thinking about the sinister use of AI or the development of AI, you could be doing it in a garage in Wimbledon. I do not wish to characterise Wimbledon as a crime centre. Nevertheless, that in turn puts a huge amount of pressure on your intelligence resources, which must be something that is moving up the agenda.

**James Cartlidge:** That is a very good question. I will bring in General Tom on this point.

**Lieutenant General Tom Copinger-Symes:** I think the point about the democratisation if you like, the proliferation, of these dual-use technologies that can be easily weaponised is profound and, of course, just emphasises the point the Minister made earlier that we need to be moving very fast on this while also being within an ethical framework. The point about advantages for non-state and rogue actors is that it is a competition, is it not? It is our job, on your behalf, to make sure that they do not get ahead, or at least they do not stay ahead if they do get ahead, and we get them back.

We have talked a little bit about recognition technologies, and it is harder to hide these days—to your point about the plutonium enrichment. We have seen in Ukraine that it is very hard to hide, but our very nature is that we will find ways to hide that spoof the things finding us. No whistle calls time on this and counts the score. It is our job, on your behalf, working with people like the Home Office, to make sure that we are constantly on the alert for how they are using dual-use technologies or anything else and that we are responding and catching up and getting ahead of them. That is the sort of immutable game that is our business. That is why we need to do all this and do it responsibly but keep doing it as fast as we possibly can to keep ahead.

**Paul Lincoln:** If you put this in the context of cyber, for example, where there are some obvious parallels here, this is not new in that sense. The UK is the third most cyber-attacked country in the world after the US and Ukraine. The National Cyber Security Centre publishes the key threats in its annual report. There are five key sets. There are four state threats—Russia, China, North Korea and Iran—but there are criminal and other threats too. The MoD, as an organisation, is also obviously targeted in that sense. We see nearly 7 million attacks against the MoD every day, so we see this kind of threshold of operation all the time and we work hard with other partners to try to counter what we see from people at large.

**The Chair:** Having been a data controller in this building, I still bear the scars. That was 10 years ago, so I can well understand the proliferation of attacks and threats.

**The Lord Bishop of Coventry:** Another theme that has been running through our deliberations is respect for the ethical principles that have been developed, even by those perhaps almost cautious or even negative about the use of anything that approaches autonomous weapons. I think that needs to be said and acknowledged. At the same time, there has been a persistent question: what does that look like in operation, and how are those principles operationalised? In that fascinating exchange between you and Lord Houghton, Minister, we saw something of the pressure on them because of the weight of responsibilities and, as you movingly put it, the equal ethical responsibility to keep people safe.

Would you be able to tell us more about how those fine ethical principles, well developed and developed with other colleagues beyond the MoD, then become operationalised? Is it a case of taking down the draft articles of the GGE to another level of manual? I do not know how you do that. Does the MoD—this sounds rather arrogant—need help to do it?

**James Cartlidge:** It is a very good question, Lord Bishop, and you heard the general talk about the training for personnel, which is the operational context. To enhance the point about the ethics of the need to maintain competitive advantage, the most moving experience I have had as a Defence Minister was when I visited Salisbury Plain to see Ukrainian civilians. In front of me was a builder, a management consultant, a couple of teachers and a student. These are, dare I say it, ordinary people who had come to Salisbury Plain to have six weeks of training and would return to trench warfare, something we have not seen in this country since D-Day preparations or possibly even the First World War, and remembering the intensity of the battle they would then face.

To me, it underlined the point that they were going out with British kit, although they would pick up further kit once they got out into Europe, and that we have a duty to stand up for our beliefs and ethics in part by defending ourselves and by supporting our allies in helping to defend them. That is by providing what are, after all, lethal weapons. It is well-documented that we have provided lethal weapons to Ukraine and there is an ethical underpinning for that, which is that they have been invaded by a tyrannical state and, if we had not done so, would almost certainly



be living in horrific circumstances under that despotic tyranny themselves. This is the balance we have to strike.

We have a powerful mission in the Ministry of Defence, which is to ensure that our military personnel can defend themselves and in turn can defend the public, the people who put us into office and who expect us to defend them. I cannot repeat enough that, at all times, it is subject to international law. I am not sure that it needs an amendment to articles or anything like that. What I think it needs is for other actors in the world to recognise the consequences of nefarious action versus countries that abide by international law, like the United Kingdom and our allies.

**Q186 Lord Clement-Jones:** Yes. I want to follow up exactly on that point, because we seem to have a kind of two-tier concept here. We have the international humanitarian law—you do your Article 36 review to see whether it conforms to international law—and then you have, in a sense, your domestic ethical principles that have been applied. I understand the context-specific aspects—the training and so on, embedding those ethics in behaviour—but we seem to have a great resistance to wanting to embed those ethical principles in a more general convention, a broader international agreement. We are keeping our ethical principles to ourselves, in a sense, which seems to me to be a rather strange policy position. Obviously, we would like others to adopt those principles.

I do not know whether you can say NATO has a common set of principles or not, but it does seem a bit strange that we have a legal framework but no international ethical framework. It looks like we will not have one any time soon, given the position of the UK at the CCW and so on.

**James Cartlidge:** Hold on a minute. First, we do have international humanitarian law. I said that its principles have been around for over 150 years. On Monday, I had the great pleasure of visiting Sweden, a country that is joining NATO because of what has happened on its doorstep. When we look at how to measure the ethical principles of a nation and its support for the concept of defence, I wear the badge. You can travel all around the country and still see Ukrainian flags hoisted. These are examples of it. They show that there is consensus and strong support, recognising that the country has been invaded and we want to support them. These are deep principles. When you measure civil society, and political institutions, there are many different ways, but ultimately there is no easy way to just conjure this up. It is long standing. We are very fortunate as a country that in our DNA we believe in freedom and the freedom of nations. That is enriched throughout our society, and of course we work with partners who share those values.

If you look at Ukraine, I am very proud of the fact that, certainly militarily, the UK rose to the occasion very promptly. It has done a huge amount to support Ukraine. It was doing so even before Ukraine was invaded, training soldiers in an operation going back to 2015. But we then used that position to cohere other nations to support and recognise. I think it has been transformative in international relations. Some countries that have been—to put it politely—laggards in some respects

have come out, to whatever extent they are capable within their own domestic circumstances, to support Ukraine as far as possible or to lend support internationally to the various motions in the UN, et cetera.

I think it is very clear where our ethics sit. We are very proud of that, and it motivates a policy that, at this very moment, is effecting those principles in supporting a country to protect itself from the invasion, which is illegal.

**Lord Clement-Jones:** I do not disagree with a word of what you have said, but you have not answered the question of how you internationalise those.

**Paul Lincoln:** To go back to your question about whether or not others are doing this and whether or not we are playing a role in that, the answer to both those questions is yes. The US and the French, for example, also have panels that they have adopted to take forward ethical principles, and NATO has a set of AI ethical principles as well. Therefore, we are working with Five Eyes partners and with other international fora, which we discussed before, not only on international humanitarian law and how it applies to artificial intelligence but on ethical principles.

**James Cartlidge:** I guess what I was trying to say—perhaps I should have been slightly more concise—is that all of that is true, but I have always taken the view that you judge people by what they do in practice. You look at the world at the moment. You look at Sweden joining NATO, which is an extraordinary change in its policy. I was in Poland, which is going up to 4% of GDP on defence. You look at the way we have all rallied around. We have towns where they have donated enormous amounts of stuff to Ukraine. Countries all around the world have done that. That in itself is an expression of ethical principles and doing the right thing. There is the framework, but there is also the practice, and that is what I was stressing.

**Lord Clement-Jones:** Nobody was saying you were not abiding by the principles. As you know, defence policy, particularly towards Ukraine, is highly cross-party. What we would like others to do is to be as ethical as us. That was the point that I was making.

**Paul Lincoln:** Beyond what I just said on the NATO perspective, we have also been, along with the Ministry of Defence and FCDO, promoting the response for the use of ethical AI approaches. There was a summit hosted by the Dutch on this in The Hague back in February this year. Again, we are trying to make sure that we are espousing the way we have adopted our ethical principles, along with other international partners.

**James Cartlidge:** To be blunt, there is so much activity we are involved in internationally. We would go beyond 12.30 pm if we went into all the details of it, honestly. I am happy to write and confirm that. There is a huge amount of activity. I think it was Lord Browne's question about Russia. There will always be challenges in that, but we are leading those

conversations, we are leading the work with allies. In all these various fora, I think we punch above our weight, but the best way to prove all that is to show how it works in practice.

**Lord Clement-Jones:** I am sure we would like to hear more about that, if you could do that.

**James Cartlidge:** I am more than happy to do so, yes.

Q187 **The Lord Bishop of Coventry:** I will take us back to the point I was trying to get to, which was less about the high-level ethical principles, which you have articulated very powerfully. It is this critique that we have heard consistently of, "Well, these are impressive ethical principles, but how are they then applied in practice?" Yes, you have listened to others on the ethics panel—that has been acknowledged—but when it comes to translating this into operational practice, are you confident that that is being done well? Do you have any sorts of systems for testing it?

**Paul Lincoln:** I hope to be reassuring. At the moment, we are developing a thing called a joint services publication. You might say, "Well, that sounds a little bit boring", but we drive through policies within the Ministry of Defence on a range of different areas through joint service publications, which we can then test against to say what we are doing. As part of that, we are requiring that developers ensure that they are embedding the AI principles within the systems. We are basically putting in there a principles-into-practice set of steps and procedures. That will require a hazard analysis to be conducted that looks at the use of AI and whether there are data failure needs and so on—going back to the life-cycle approach that we talked about before. Then we will put a process of verification in place around that.

If we put this in the context of the overall governance that we have in the department on this, there is a set of policies that sits in the head office. The delivery into military capability is done through our normal military capability, but the instructions coming out and doing that will contain some of these components, and the coherence across that is tested through the Defence Artificial Intelligence Centre, which is part of Strategic Command, which is where we do the standard process.

**The Lord Bishop of Coventry:** That is very helpful. A little question in my mind is whether any sort of case could be made for further involvement in the ethics panel or some other such body in analysing, as a matter of principle, whether there is the sort of coherence that you talk about. Of course, there will be all sorts of constraints there.

**Paul Lincoln:** I am sure we can reflect on that, Lord Bishop.

Q188 **Lord Clement-Jones:** Can we hear a little bit more about the ethics advisory panel and the frequency of meetings, their composition and so on? It depends on what time we have left, but it would be very interesting to hear what processes it has, what its membership is, how often it meets and how issues come to it.

**Paul Lincoln:** It has met five times. We set out again in the document that we published on the “ambitious, safe, responsible” approach to AI that we established this committee. This was a proactive thing that we did. It started in March 2021. The sixth meeting is next week. It has gone through a range of different things, including quite heavily at the start how we will develop and adopt the ethical principles. The membership, which you touched on as part of that, again is being updated, because people have moved posts and all the rest of it, but, broadly, the membership again was published in that document. It involves combinations of defence, industry, academia and critics of government policy as well in this space to make sure that we have a balanced view that is being taken into those considerations. The principles that we came up with also take account of the wider government Centre for Data Ethics and Innovation and the approaches that it has been taking more widely across government.

**The Chair:** You will know that Professor Mariarosaria Taddeo has already given evidence to this committee, so we have had an insight into that. I think it would be helpful if those points could be followed up in writing and we could have the details of those.

**James Cartlidge:** Just to clarify, this is on the work of the panel and a sort of list of the many engagements we have on the international front.

**Lord Clement-Jones:** That would be terrific if you could do so.

**The Chair:** On the panel, as it were, the work programme, frequency of meetings, membership and so on, which we have, I think, in various forms, but it would be helpful to have it all together.

Q189 **Lord Fairfax of Cameron:** I am just trying to apply some of the things we have been discussing in the last 10 minutes to a real-life theoretical example. Although it is theoretical, it may become a very realistic one very soon. What if this country was being attacked by a mass drone swarm—100,000 mini-drones; it does not matter how many—and, in response, we had our own defensive drone swarm, so unmanned against unmanned, and a human here then authorised the release of a defensive drone swarm to confront the offensive drone swarm coming from the aggressor?

I am assuming that, AI-enabled, our defensive drone swarms would then be able to make decisions for themselves about how best to engage the offensive drone swarm. Would that all be fine? I do not see that offending against international humanitarian law—it is unmanned against unmanned, for example—but would it be okay, because the context-appropriate human involvement requirement had been satisfied by the original human command to release our defensive drone swarm? I think it is not too complicated.

**James Cartlidge:** A couple of points. This is very much a hypothetical discussion, and it would be an operational decision as to how to respond to any specific threat, which would be according to the circumstances.

**Lord Fairfax of Cameron:** We only took that one, because it is one that some people think might become quite realistic before too long.

**James Cartlidge:** The Second Permanent Secretary is itching to come in on it. I will pass the ball shortly. On the second point—Paul made this point about the naval gun, Phalanx—you have to think of that context-appropriate human involvement as a life cycle. The programming of those capabilities would be such that the parameters will be drawn very tightly. Even if they were able to swerve or do whatever kind of activity—as is arguably the case, for example, once some of our missiles are released—as long as it was grounded in those parameters and was not fully autonomous, arguably it would therefore be compatible. This is an entirely theoretical discussion.

**Lord Fairfax of Cameron:** In my example, our mini-drones, defensive ones, are fully autonomous, but they can be released in the first place only by a human. I am talking about their becoming fully autonomous once released. That is my exam question.

**Paul Lincoln:** I would still, as the Minister said, bring that back to the whole life cycle and the parameters associated with the design of those mini-drones. Perhaps a more realistic here-and-how example of where you might say there are systems that would be useful is what has been going on in Ukraine, where a US-supplied Patriot missile system is taking out incoming attacks from Russia. Because of how those are programmed, they need to be able to take out systems that a human would not be able to do otherwise, with speed and precision.

**James Cartlidge:** A point of information, Chair, which I thought might be helpful. On your question, Lord Bishop, I asked if there were any clergy on the advisory panel, and I am informed that Professor Peter Lee, professor of applied ethics at the University of Portsmouth, was an RAF chaplain from 2001 to 2008. I thought that might interest you.

**The Chair:** Thank you very much.

Q190 **Lord Clement-Jones:** We have heard you use the phrase "the public" quite a bit during the discussion, particularly in the context of where to draw the line and drawing the line at fully autonomous weapons and so on. How do you know what the public will accept, and how do you measure public support for the current policies? Do you consider that you have, for instance, democratic support for current policies? How are these being discussed, in what forum, and how will you be confident that the public is behind whatever policy we adopt going forward?

**James Cartlidge:** I think one thing that will not change, whatever happens with artificial intelligence, is that the way these matters are determined, which is in an election. That is how we determine the public support something.

In terms of our consent to this, it is important to stress—this will be true for any Government of any colour—that there is the tacit point that you are there to deliver the national interest, but you cannot have a

manifesto that covers every potential development when you are in office that you have to respond to. There is that simple principle that you govern in the national interest. I think we have shown leadership on this. The Prime Minister's summit, which I referred to earlier, albeit that it does not include defence, is a very important statement of our intent, as will the AI safety summit, which will be a global event in the autumn.

Interestingly, President Biden has expressed some of the anxiety on this. He is the most powerful directly elected leader in the world. The point about this sense of anxiety is that it is, I think, true of technology in many ways. I pick it up from my electorate in south Suffolk. That is why one of the most important points is to have sessions like this, where we can be as transparent with people as possible about the inevitable risk of non-state actors, for example, which is perhaps one of the key risks, but also about the opportunity, which could be manifold. For defence, I think much of that will be not about the sorts of applications that we have been talking about, but about data analysis and safety, such as defusing bombs and so on.

Ultimately, it about us—a key point I made to the Lord Bishop—maintaining our competitive edge by investing, interrogating this technology and ensuring that we have projects that mean we are at the maximum fighting edge. I think the public support that. They would expect that of any Government of any colour. These are basic principles of how we govern in the UK. That does not mean that we would have a referendum on it, and I do not know what the specific opinion polling is, but I think it is about reassuring the public and showing leadership. I am confident that, as a Government, we are doing that.

**Lord Clement-Jones:** That begs a question, though. You talk about a global event, the global AI safety summit, and son on, which will be splendid. We applaud that initiative, but of course that omits the fact that there could be a very good case—and a couple of our witnesses have certainly made that point—for more of a domestic discussion about defence AI. I entirely agree with you that public anxiety is not purely about defence AI systems—far from it. The existential risk language is used in all sorts of other contexts as well, but would there not be a case for a domestic conversation about this to a much greater extent?

**James Cartlidge:** DSIT is, of course, the department that leads on AI across government. It has published its strategy. We have our own MoD defence AI strategy; we have a copy of it here. This is part of that engagement. At the end of the day, when you are in government you have to lead, you have to show that you are ahead of events and ahead of technology. I think we are doing that as far as we are able. These are complex matters. The potential range of threats is huge. We have been alluding to the “what might happens”—dare I say it, Donald Rumsfeld’s unknown unknowns—which in AI could be quite extraordinary.

The point is that we have set out a good foundation in our department. Paul has spoken about the various frameworks for skills and getting people in the right place. As a department, we are confident that we are

doing the right thing, and as a Government we will be showing leadership through that summit, but through DSIT, through working across government, I think we have the right blend of policies to ensure that we can both balance the risks that may be present in AI but also bring the public with us in seeing the opportunities, not least economically, and, from our department's point of view in terms of innovation, going forwards in our capability.

**The Chair:** Thank you very much. Minister, I am not sure I will let you get away with the assertion or the implication that a general election result provides authority for the range of government policies that may follow. I think perhaps dealing with those individually and scrutinising them does fall to Parliament. Equally, if you want to seek greater public endorsement and approval for AI and the associated systems, an understanding will be enormously important.

**James Cartlidge:** Chair, you are right, but I want to clarify. All I was saying is that we draw our authority from the election result. That is how our system works, that is the legitimate authority. But, as you know, things then happen. Who knew that there would be a pandemic? Things happen, and it is about tacit support from the public. They expect you to do what is in the national interest.

**The Chair:** Oh, events, dear boy, events, absolutely, but I hope that if you are seeking greater public endorsement and approval, you will agree—I am sure we would all agree—that a greater understanding of what is involved is also essential. I think my colleagues and I hope that our inquiry will play a part in that.

**James Cartlidge:** It is very important in that regard. I have a NATO definition of autonomous. I would be happy to write to you with that or read it out. It is entirely up to you. It is quite long, so perhaps we will write.

**The Chair:** We have covered a very broad canvas this morning and we have quite a few points on which you have very kindly agreed to write to us, which we will confirm to you. In the meantime, thank you very much from all of us for giving your time and expertise. We are extremely grateful.

**James Cartlidge:** You are welcome. It has been a real pleasure. Thank you very much. We will write to you, as discussed.