



Artificial Intelligence in Weapons Systems Committee

Corrected oral evidence: Artificial intelligence in weapons systems

Thursday 11 May 2023

10 am

[Watch the meeting](#)

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

Evidence Session No. 5

Heard in Public

Questions 64 – 80

Witness

I: Professor Sir Lawrence Freedman, Emeritus Professor of War Studies, King's College London.

USE OF THE TRANSCRIPT

1. This is a corrected transcript of evidence taken in public and webcast on 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 witness

The Rt Hon Sir Lawrence Freedman, KCMG, CBE, FBA.

Q64 **The Chair:** Professor Freedman, it has been many years since you and I last collaborated on defence matters, but it makes it all the more welcome to have you at the committee this morning. You are a very old hand at this, and you will know this is being webcast, and there will be a transcript where you can correct any factual inaccuracies. I hope that we can finish by about 11.30 am, if that suits everybody.

I wonder if I could begin by picking up something you said in a previous publication about how technological developments encourage a fantasy as to how warfare can become “fast, easy and decisive”. We are moving on in technological capability at an extraordinary rate, so how does that remark now play against the background of AWS? Are there fantasies, as it were, which we need to explode? That may be the wrong word, but you understand.

Professor Sir Lawrence Freedman: I understand. Thank you for having me here; it is always an honour. My concern is with two things. The first is technological determinism: that new technologies mandate a certain form of warfare and that war follows where technology leads. I will explain in a second why I do not think that works. Secondly, there is a more general problem—you can call it a fantasy; it has disastrous consequences—which is the belief in the decisive first blow, the surprise attack that solves every problem. Technology is often seen as a means by which you can achieve things that you were not able to achieve by other methods.

You have to step back a bit and consider generally how war has developed. It is layered. We can see this dramatically at the moment in Ukraine. You can watch videos on social media of troops engaged in almost hand-to-hand fighting in trenches, something that would not have surprised anybody had they had such imagery from the First World War. At the same time, you can get videos of Ukrainian troops using their iPads to connect to Starlink to bring in intelligence coming from satellites, and using another satellite communication to direct artillery accurately to hit particular targets. These things happen together.

So it is not that technology is irrelevant or gets lost in the timeless character of war, but it is layered. It comes on top of what is already there, which normally takes up the bulk of activity, even if you just call it legacy systems in terms of equipment.

The second point, about the problem of the decisive first blow, goes right back to the 19th century and the belief that war is decided by a battle. If you win the battle quickly, you will be in a commanding position to get the political settlement you want, and because wars are decided by battle the best way to win is to achieve surprise and catch the enemy unawares, and this will make all the difference.

A historian of these things can find these themes re-emerging all the time. The technology that made a difference in the late 19th century was the railroad, because it made an enormous difference to logistics as well as the gradual improvements to artillery, which were more incremental. In some of the fantastical, fictional works about an invasion of Britain that appeared in the 1870s, the first thing that was done was cutting off the telegraph line. These issues have always been around.

The technological dynamic aspect of it became more important in some ways as the prospect of battle receded. This leads us directly into AI. With the first moves in the digital age, particularly the development of cyberwarfare, you will see references to an electronic Pearl Harbour, that one day the plug will be pulled out, all our systems will collapse and everything will go black because of some massive cyberattack. One thing to think about is that this was exactly what was expected on 24 February last year. It was not that the Russians did not attempt a cyberattack but that they failed, and they failed because people knew that it was coming and had taken measures to mitigate the effects.

These are my concerns. It is not that we should dismiss the technology or play down its impact, but that we must put it into the wider context of why and remember that war is a duel; it is not one side just imposing its will, except on occasion, but two wills against each other, and they resist and find ways to adapt.

Q65 The Chair: We all know the political attraction of nice new shiny systems that promise so much more than what went before. What sort of warnings, political directions, do you think are appropriate, apart from reminding politicians about this layered nature of warfare?

Professor Sir Lawrence Freedman: It is always useful to do that. One of the difficulties with AI is that it does not appear as bright and shiny. It appears as remarkable outcomes, but it is not like a precision-guided munition, for example, which did appear bright and shiny when you could guarantee that if you got the co-ordinates and the range right, you could attack and be pretty sure about destroying anything. If you go back to 1991 and the Gulf War, that created a whole set of new images that made an enormous impact, because you could see it happening.

One of the challenges of AI is that we see our smartphones being able to do remarkable things—even our students are discovering how somebody else can write their essays—but it is not bright and shiny. It is a continuation of what has gone before but with a cleverness that people find difficult to understand. That is why people, including me, find it quite hard to get their heads round the full implications of what is going on with AI.

The Chair: Lord Mitchell has a supplementary question on this.

Q66 Lord Mitchell: I do. Hopefully, Chair, you will slap me over the wrist if I go off-piste. You mentioned Starlink, Professor Freedman. Just to refresh, that is Elon Musk's low-orbit satellite.

Professor Sir Lawrence Freedman: Yes.

Lord Mitchell: One of the things I have read about it is that he is deciding, on Ukraine, who can access it and who cannot. Here we have a situation where a high-tech businessman is deciding foreign policy, which has implications for all the things we are talking about. Do you have any comment on that?

Professor Sir Lawrence Freedman: You are right to raise it; it is an important issue. The relevant technologies that are affecting so many human affairs, not just warfare, have, since the 1990s, been developed by big companies, largely, but not solely, in America, and they have an enormous impacts, so it must come in with Starlink. At the start of the war there was the Viasat satellite, which the Ukrainians had been using for communications. The Russians basically took that out by jamming it very effectively right at the start. Then Musk came in—if he had not, there would have been problems—with Starlink providing thousands of terminals that enabled Ukrainian Government communications to continue, allowed Zelensky to keep on talking to people, and helped at the front line.

However, as you say, Musk developed his own foreign policy; he literally proposed his own peace plan. When it looked like the Ukrainians might be able to attack into Russian territory, he put a limit on what they could do. It was not so much who could use it; it is just there were limits on their ability to use it to direct stuff into Russian territory.

You could also look at the role that Microsoft has played, which has been absolutely vital in Ukrainian cyber defences. Microsoft, without getting the same attention as Musk, has played a very important role. Again, one of the things we need to get our heads around is that, although there are some forms of AI that only states can and will develop, an awful lot is being developed outside of states. ChatGPT, and so on, is not being developed by states. China is slightly different in that, but on the western side it means that the role of corporations has become important. You may recall after the Snowden revelations—we are now talking about well over a decade ago—there were deep anxieties amongst many in this community, Google in particular. They did not want anything to do with the military. They wanted technology to be developed in a very benign way to address problems of climate change and suchlike, but not problems of defence.

Interestingly, one of the consequences of Russian aggression is that people can see the point of defence a bit more now and understand why it might be necessary for Google and others to take contracts and support the Pentagon in developing their AI work. This business/military interaction is really important. As you say, if you have a corporate leader like Musk—I have to say that he is a bit unique in this—or somebody like that with their own agenda, this is somebody else who has to be negotiated with. At one point, it looked like Musk was going to pull the plug on support altogether, because it was costing him money and it was expensive. Quite complex negotiations had to take place just to keep him

doing what he had been doing up to that point, never mind going the next step.

Q67 Lord Browne of Ladyton: Good morning, Professor Freedman, I am pleased to see you here. This question is designed to allow you to share your reflections on the UK's policy on AWS, and the extent to which the UK Government aspire to be a leading voice in conversations on international regulation. I would like you to address a couple of things. One is an observation. Another is a very direct question.

The evidence that we have received up to now is hardly surprising, because it confirms that we are aligning ourselves very closely with the United States, although we are not aligning ourselves as closely with it on transparency about its own development of weapons systems. In the context of the CCW, that has meant expressly that we have aligned ourselves with Australia, Canada, Japan and South Korea, which are a very close small club of nation states and not necessarily representative of the international community. What might an arms control regime to regulate autonomous weapons actually look like in the current international context?

Professor Sir Lawrence Freedman: On your first point, you are clearly right. The UK is, by and large, on the light-regulation side of the debate—lighter than the EU, for example. AI is an important feature of the second pillar of AUKUS, with Australia and the United States, which, in principle, is an area that requires almost an alignment of regulations, so you could have the technology sharing which that envisages, which over the longer term could be really quite important.

The question of an arms control regime is separate. I cannot say that any arms control looks very promising at the moment. It is worth noting all the enthusiasm for demining and avoiding mines in warfare, and seeing just how many mines are currently being laid in Ukraine. The imperatives of war take you in directions that go against the demands that tend to lie behind arms control aspirations.

It is also very difficult to think through what it is you are trying to stop. Autonomous vehicles, for example, first received a lot of attention in warfare as a means of dealing with IEDs—improvised explosive devices. It made an awful lot more sense to send a machine to probe and prod something that looked suspicious than to send a human being. You would not want to stop that. It is a good thing to do; it is better than having humans in that position.

If you start to imagine robots, or swarms of drones, going into war, you have another problem, but the technology that allows you to start directing drones to do reconnaissance may not be that different from the technology that allows you to direct them to drop grenades on trenches. Again, in the current war—this is true on the Russian side as well, but we know more about the Ukrainian side—a lot of improvisation goes on. It almost goes back to the previous question and the times when everything was under the control of the state. If the state had signed an agreement

not to do something, you had some trust that it might not happen. Actually, it is much more fluid now. The participants in conflict are much wider, as is their ability to direct technology itself.

Quite a lot of work has been done on the possibility of arms control regimes, but it was done in slightly easier times, and I have to say that, at the moment, I would not be optimistic. The restraints will have to be self-generated. We will have to work out the things that we do not want to do and would rather not see, and hope that others are thinking in similar ways.

Lord Browne of Ladyton: If I may bring you back to the original question, there is an international process, which has been going on for a long time and has not produced very much. The question is: are we leading in that, or, as many may think, are we a brake in that process?

Professor Sir Lawrence Freedman: I do not know enough to give a definitive answer. I think the steam has gone out of a lot of that work at the moment. One of the difficulties with international negotiations and this sort of technology is that you can find yourself addressing a problem that is a year and a half, or even five months or a month, out of date, because something happens that you were just not anticipating. You can work on general principles—the importance of keeping humans in the loop, the dangers of misinformation and fakery, and all these things—but what I find difficult to see is how this can become the basis of a reliable international agreement. It is important that we think about the ethical and legal issues connected with AI. There are lots of groups doing a lot of work on that, so it is not as though this is neglected, but it tends to be on a national or European-wide regulatory basis, or maybe within the alliance, but you will bring in China only with difficulty, let alone Russia.

Q68 **Lord Houghton of Richmond:** The formal supplementary question in this set is whether existing battlefield regulations, including rules of engagement, targeting directives, and command and control procedures, are sufficient, given the use of AWS. We want to get your views about the extent to which what is called meaningful human control should be a safety mechanism within the potential fielding of artificial intelligence in weapons systems.

These sessions are a bit like staff college syndicates, where half the question is about impressing the directing staff and the other half is about trying to influence the views of your compatriots, in the old military manner. In the sessions so far, I have tried to let the others know that, to my perspective, the battlefield, although hugely chaotic, is a highly regulated environment. It has all these things like targeting directives and rules of engagement—weapons tight, weapons free, stand to, all those things. I have not mentioned an OODA loop to them. OODA comes from some ex-pilot from America.

Professor Sir Lawrence Freedman: It was John Boyd.

Lord Houghton of Richmond: It stands for: observe, orient, decide,

act. It is one of the rhythms of battle that you are continuously doing. You observe, you orient, you gain your information, that leads to a decision, and then you act. For weapons systems, we could probably reduce that to: inform, decide, act. Every element of that inform, decide, act loop could be subject to and linked by AI.

There are some systems fielded today, particularly those in the self-defence of surface ships from inbound air attack, where the whole thing happens within a context of speed and response, and relative ease of discrimination and proportionality, which you could delegate in a fully automated sense. But you could also have meaningful human involvement in every stage, or just one or two stages, of that, depending on the varying contextual considerations of time, discrimination, proportionality and legality.

I see no difficulty in the ability to do that, because I come from the fundamental premise that, as a battlefield commander, you cannot ultimately delegate legal or moral responsibility to a machine, but you can exercise judgment in what you might call permitting automaticity. I can get my head around that, but it involves an awful lot of high-quality training in an understanding of the variable geometry of the permissions that you would need to take a risk judgment against in order permitting automaticity.

I think this is all manageable within what we have, but I would be hugely desirous that you challenge that, or else help me to win my fellow colleagues on the panel over to my view.

The Chair: You see a table full of open minds here.

Professor Sir Lawrence Freedman: It is just like a staff college syndicate. You could argue that the problem of working out whether something is a legitimate target may be exactly the sort of thing where AI can help. You may recall the endless arguments that took place over targeting during the Kosovo war. Some of these were political, such as how far you could go, and what was and what was not appropriate. Some of them required particular information or intelligence about how things had been changing over time. Then you have to bring in the international humanitarian law and so on. This is exactly the sort of thing where AI does not have to make the decision for you but can give you the material very quickly. That is where it can be useful.

On the question of what to do if you have to activate your defences, again, this is exactly where AI makes an impact where it is in use and you can see it being developed. It is a good way to get into how information technology has made a difference over decades. You will recall the Dowding system in 1940 and the Battle of Britain, which required a combination of radar and human observers sending information to Uxbridge and putting all this information together. They knew which RAF wings were available, and they had to work out what was most important, where they were going, and what was available to

deal with them. They dealt with it very quickly, but it was all done by human beings moving those things around on the table.

Fast-forward to Israel's Iron Dome. It is doing exactly the same thing; it is working out where incoming missiles are likely to land, which ones can be safely ignored, which ones must be dealt with, what is available to deal with them, how you prioritise, and it does that in seconds because that is all the time they have. There is just no room for human intervention in that process. It works because it is defensive, because people understand that if they had to have a human being deciding all these things it would just be too late.

That is true on a ship. Going back to the 1991 war, there was a famous episode where an officer on a Royal Navy ship had to decide whether to intercept a missile that was going towards an American ship. The risk was that it was a civilian airline, not a missile. He decided it was a missile, and correctly attacked it. They could not work out how he knew. Eventually, they went back and saw that the radar trace began at the shore. That obviously meant that it was a missile rather than an aircraft that had come from some distance.

This is normally taken as an example of intuitive thinking, of how you can make good decisions even if you are not deliberating because you do not have time. All that could be worked out. That, again, is the sort of thing that you can assume the new technologies will facilitate. When you are activating defensive systems, that is all you may be able to do. It will be very different, though—this comes back to the first point about how you decide how to target—when you are working out what to attack and why, when you are taking the initiative. That is where, it seems to me, you absolutely have to keep humans in the loop, because you do not want to take these bold decisions on the say-so of a computer, however good the algorithms.

Lord Houghton of Richmond: On that last point, the danger appears to be if the person who is delegating automaticity to the machine does not genuinely understand the potential envelope of lethality of what it might do—in other words, that it somehow becomes sentient and goes off on its own. So long as a weapons system can be so tested and proven, that the individual commander making the delegation of automaticity knows the envelope of lethality he is unleashing, the humanitarian laws of distinction and proportionality can still be adhered to. Is that a fair comment?

Professor Sir Lawrence Freedman: Yes, it is. In war, things go wrong. Systems that you thought were wonderful can break down. Other things land in the wrong place and cause enormous and terrible damage that nobody expected or wanted. This is why war is a terrible thing. At the same time, the more narrowly you can frame the problem and the more specific you can make it, the more likely you are to be able to stay in control of how things develop. That is my general rule for trying to understand the role of AI. It can deal with narrow problems where there is a lot of information, but the less reliable the information, the more it

may lead you astray. The more widely you set the problem and the more you ask it to do, the more confusing it will get. That is one way, it seems to me, of keeping it in control, which is why the air defence problem is a particularly good example.

The Chair: I should observe that Philip Wilcocks, the captain of HMS Gloucester, was a very good friend of mine. Alas, I attended his funeral last month—much too soon, much too early.

- Q69 **Lord Hamilton of Epsom:** I would like to go back your concerns, Professor Freedman, about international law and whether it has any role to play in this. You are not the first witness to cast doubts on that. If we are to look at it in terms of the United Kingdom, is there a national solution to regulation where there is no international solution? Could some committee be set up by Parliament that oversaw the procurement of AI and, indeed, its use on the rare occasions when we went to war?

Professor Sir Lawrence Freedman: We have international humanitarian law, and, as I understand it, a lot of the discussions have been about whether you need anything extra or whether that can cope. I am not wholly sceptical of international law, but it tends to work best as broad principles rather than trying to itemise everything that can and cannot be done. The principles that you see under development and being discussed for AI in the civilian sphere tend to follow that.

The problem is interpretation and application. I find it very difficult to say whether AI is the sort of thing that needs a committee by itself or whether inquiries such as this can deal with it as part of a broader framework. It is a problem of very fast-moving technology. I suspect, and I am getting well out of my comfort zone here, that it will also be a question of case law: what challenges have come up, and how have they been addressed in the past? I am a historian, as you will see. I keep going back to past wars, and I can think of examples that illustrate particular things, but we are in a very new area now and we may not have the examples that we were anticipating.

Finally, we keep coming back in this area to something that you may want to discuss a bit more: that it is not one state imposing its will on everybody else. We do not expect that. It is states engaged with each other and trying to frustrate each other, which makes it very hard to work out exactly where the boundaries should be and where the limits are, because you do not quite understand what sort of resilience you need. It is very difficult to develop very specific instruments to deal with these developing problems. Principles, yes. Applying international humanitarian law, certainly. But I am not so sure that we will find it very easy to take it beyond that.

- Q70 **Lord Hamilton of Epsom:** That half answers the next question. Can we create a defence policy that is resilient to the rapid change that is going on?

Professor Sir Lawrence Freedman: I think we can. It comes back to my opening point that we have to remember everything else that is going

on. It is only one part of the defence policy. We have been living with the idea for over 30 years—40 years, in some ways—that modern warfare is all about smart bombs and smart weapons and so on, but you can run out of smart bombs and smart weapons very quickly, as we saw. You go back to dumb bombs quite quickly. The methods by which Russia is fighting at the moment rely an awful lot on artillery barrages and rather hopeless infantry assaults, so they have gone backwards as they have used up stuff.

As you will know, one of the challenges for British defence policy at the moment is filling an empty cupboard. So much stuff has had to be passed on to Ukraine—correctly, in my view—that there is not a lot left. If you do not have your basic ammunition, there will be trouble picking up everything else. That is why you have to see the impact of new technologies in very particular areas. It is not going to transform every aspect of war. It can mean that an infantry company has better situational awareness than it had before, but if it does not have the artillery ammunition to fire, that will not do it much good.

The challenge is to keep these different layers integrated and not to neglect the old ones, because in the end it is the weapons that can make all the difference in war. The challenge is to identify the areas in which AI can really make a difference. Remember, it is basically about decision-making. AI itself does not kill people or destroy things. It supports decision-making that enables other things to do that. Obviously if you link it with autonomous weapons and so on, it starts to give you a capability that causes you to wonder about the role of humans.

I am sorry to keep going back to the current war, but it is the best evidence we have. If you look at the way drones have developed, even during the last 14 months, and the roles they are playing, you see some of the more sophisticated drones becoming quite problematic because they are jammed because electronic warfare techniques can be improved to deal with them, which is why you do not hear about the Turkish drones so much. A lot of other drones—cheap commercial ones, many from China—are being used on the front lines and adapted in some ways, often quite cleverly. It does not matter if a lot of them are lost because of electronic warfare. You may think that it is a quite good deal if they are knocked down by an air defence missile.

It has not gone in the direction people anticipated. I am very careful about drawing lessons from a war that is not yet over, but, at the stage we are in at the moment, drones are becoming more important but cheaper and less sophisticated than some of the ones with which we started. That is linked with Starlink and the intelligence feeds and so on. It is not as though that element is absent, but sometimes a robust, basic technology does you just fine. AI makes a difference in the decision-making, providing you with the evidence and the options. The risks are there, but so are the possibilities.

Q71 Baroness Hodgson of Abinger: Can I pick up very briefly on what you are saying about the decision-making, which I absolutely understand. I

can see that AI is very good at systematic decision-making, but it lacks empathy. Could that lack of empathy being fed into complex decision-making be problem?

Professor Sir Lawrence Freedman: Indeed. It is not going to give inspiring speeches or pat somebody on the back and write consoling letters back to the bereaved family. It obviously has none of that: it is algorithms. What it can do, on the basis of the information and the data it has, is interrogate that very quickly. If you have asked it the right questions with the right parameters, it will give you a helpful answer.

My worry, and you can see this already in lots of the applications in everyday life, is that it comes out with answers that you do not understand. You cannot quite work out where they come from because you do not know the weighting that has been given to particular factors. The first thing any of us ever learned about computers and these systems was rubbish in, rubbish out. By and large, when you were getting rubbish out you could go back and see the rubbish that had gone in that biased the decision. If you do not understand all the inputs, it can become quite bewildering.

This is a general problem. A lot of the systems that may be used now in social security or tax to test particular things may say at the other end whether you can get a benefit or not, but unless you know how well they have understood the case law and the benefit system, and how well they have applied particular questions, you cannot be sure, but that is all you have and people will be encouraged to rely on all this.

This is where the awareness of the limits of what you have becomes very important. How you get around that when lots of decisions, which may be quite low-level decisions but can be quite important, become an awful lot easier is quite a big challenge. That is the first thing. Classic examples tend to be given about decisions that machines are very bad at taking with autonomy. There is the standard driverless car problem, for example: if you are about to crash and there are pedestrians on the pavement, which way do you turn? It is very hard to train a machine to make that decision for you. All one can do is put the problem on the table and encourage those who are designing the systems to think very hard, asking the legal and ethical questions as they go along and trying to imagine the various contingencies they may face. Those are the sort of things that need to be done.

It goes back to the point I made about these systems working best with a narrow problem and a lot of information. You might trust driverless cars on a motorway but not in the middle of London, because the contingencies become too complex. One problem with a lot of AI, which is obviously eased the more complex and sophisticated they become, is that we may notice something a bit anomalous that excites our curiosity but it may be seen as a bit of outlying data which they do not understand and can ignore or it may throw them completely off course.

These are the sorts of things that are a danger: when people are trying to make decisions relying on the support of AI and do not understand how the algorithms are working, what the critical data has been and how it has been weighted. It is not just that the decisions are not very empathetic; they are also pretty poor.

Q72 Lord Clement-Jones: What you said earlier about drone development feeds very nicely into my question. There is a widespread narrative about an AI arms race, but very few examples of autonomous weapons systems exist. What is the reason for this seeming mismatch? I have a couple of examples. Dr Dear told us when he was asked to provide an example of a major AI project in the MoD that he was stumped. We have had other examples of weaponry that could be the subject of an arms race: Phalanx CIWS, apparently, and loitering munitions. I wondered what you thought about this whole question of an arms race. It appears that we may not be moving into a very sophisticated era; it may be more, as you say, about small drones that are not that sophisticated.

Professor Sir Lawrence Freedman: One has to be a bit careful about drawing lessons from what the Ukrainians have been doing, because they are not the Americans. The Americans would have fought this war very differently, because they have air power and quite a lot of stuff that is much more sophisticated than has been provided to the Ukrainians. It depends on who is fighting and over what.

The problem with the notion of an arms race is that it assumes that everybody is trying to do the same thing. One advantage the US demonstrated with its drones was loitering. In the war on terror, drones were being operated from thousands of miles away over critical targets, watching what was going on and deciding whether to unleash a weapon. Everybody can try to develop drones like that. One reason why the Americans could do that was that nobody was shooting back at them.

There are two sorts of arms races. There is the arms race where both sides are trying to do the same thing. Then there is the arms race where one side is trying to work out how to neutralise the capability of the other. That is the most important arms race, that is the one that goes on daily. You can see it in the cyber domain away from the military sphere; there is a constant duel going on as hackers try to develop new ways of getting into other people's networks, and those in charge of the networks try to find new ways of stopping them and detecting it. This is constant; it goes on all the time.

This, again, is relevant to the Ukraine war, because Ukraine has had enormous support in dealing with Russian cyberattacks, despite the fact that these have probably been the most sustained attacks ever. That is the sort of arms race that we are a part of, and AI is just a development of that. The thing to watch is not only how we deploy and support loitering drones and make the most of the information they are sending back, but what others will be doing to interfere with them and to stop them doing those things. That is the arms race, and both sides are fighting it in both respects.

All these different bits are linked, but they can be treated separately. On one hand you have the basic resilience of your networks; your ability to keep the feeds coming in from satellites. What may also be critical is your terminals being jammed and you suddenly going blind. However good your machine learning is, if they have nothing to learn with, it will not help you very much. On the other side are the weapons they are trying to direct, which all have their own histories and technologies that need to be understood. How the systems integration works, how they are brought together, is one of the things that AI is supposed to do. Different parts of this can be disconnected, as we are discussing with drones, and still make an impact, but they may have more impact if you can integrate them much more with a wider network.

That is a long-winded way of saying that it is not particularly useful just to think about everybody charging off in the same direction and seeing who gets there first. It is a much more complex interaction with a variety of component parts, and the fundamental problem of offence and defence—of the interaction between the two opposing sides.

Q73 Lord Clement-Jones: That is really helpful. We talked earlier about Starlink and so on. Does the heavy involvement of the corporate sector in AI development and so on make a bigger difference in this area than in more conventional weaponry?

Professor Sir Lawrence Freedman: Yes, absolutely. This has been a feature of the digital age. If you go back into the history of computers, they would not have developed without military contracts. With most of the mainframes of the 1950s it was a few large corporations and basically the state. As microchips became smaller and smaller and more could be done with them, Moore's law and all that took effect.

Then, in the 1970s, you had the personal computer and, all of a sudden, stuff that only states could do individuals could now do, never mind companies. The Gulf War was a turning point. Up to 1991, only the military had global positioning. Lots of reporters following the Americans around suddenly realised what geolocation could do for you, and before you knew where you were a smartphone had all these things. All the things on a smartphone now were very privileged for the military a few decades ago, but now we expect them and take them for granted.

One of the consequences of that was the civilianisation of technology. Technology is being developed by civilians; it is not being developed solely by the military anymore. Ordnance is a different matter, because you do not want to hand over too much. Companies can fulfil orders from the Ministry of Defence, but the Ministry of Defence will still give the orders; you do not want private companies ordering their own shells.

There is no sharp civil/military divide in this whole issue, and what once started with the military in the strongest position has moved to a situation where the civilian development is strongest and most important and will continue to be so. Often, we are now talking about the military catching up, so one debate that has been going on for a long time is how

much the military really need to develop dedicated systems of their own when off-the-shelf may be a lot cheaper and a lot easier, although potentially more vulnerable to interference.

These are the sorts of debates, but the bias now is certainly the big companies developing the technology, and they are international companies. Apple may still be based in the US, but if it is worried about its Chinese market, that affects its attitude to encryption and so on.

The Chair: On the relationship between defence and offence that you were talking about a moment or two ago, I seem to recall that WS Gilbert said, "Dear stalking would be a very fine sport if only the deer had guns", which may provide a sidelight on that.

Q74 **The Lord Bishop of Coventry:** Greetings from Kenya. Thank you, Professor Freedman, for your presence. I am sorry not to be with you in person. May I take you back to a couple of things you said earlier? When you were talking about a defence system, you said there was simply no room for human interaction because of the speed. Later you talked reassuringly about the function of AI being to support human decision-making.

My question relates to the warning we have been given by a number of people about the risk of escalation through speed and the difficulty of predicting the response of other actors. If those risks are real, how can that risk of increased escalation be mitigated and responded to?

Professor Sir Lawrence Freedman: Escalation is a strange concept. It can mean two different things. First, it describes a tragedy—you step on to the escalator and you cannot get off, which is the original meaning and why the term came into use—and suggests a situation in which the human beings have almost lost control because of the passion, intensity and confusion of war. In response to that view of escalation, it became a deliberate process, something that you controlled.

You may recall—it keeps coming up—the famous escalation ladder developed by Herman Kahn in the 1960s. It had 44 different rungs, and you could decide which one to go up and when. Alarming, the first nuclear one was about rung 13 and it had lots of different ways of using nuclear weapons. His point was that these were deliberate decisions. Only over time did you get out of control. To start with you were in control.

These ideas go side by side, but the first point to emphasise is that some escalation is very deliberate. You do it because you are losing and because you want to move to a new level which you think you might win. That is a lot of the concern about escalation in this war on the Russian side. What happens if Crimea is threatened? What happens if Russian forces are being humiliated? Would he go nuclear? I do not think he would, but you can see where the concern comes from.

The first sort is escalation that is inadvertent. It comes from friction—the fog of war, if you like. People are concerned that either you have general confusion—as I said before, you do not quite understand why things are

happening, why decisions are being taken by the machine and the evidence they are giving you—or there is deliberate interference. Look at the ability to create fake news—very credible fakery now—or to start a rumour.

These sorts of things can lead to politicians under a lot of pressure to make a deliberate decision but on the basis of false information that has been deliberately inserted. There was quite a lot of concern about this before the war. How do you work out what exactly has happened? As anybody who is following events closely knows, you hear a report of some atrocity, which is immediately denied, and it takes a while to work out exactly what has been going on. Sometimes it turns out to be true, in which case that has an effect on the readiness of the other side to escalate.

You cannot just say that escalation is a negative. If a country is trying to win a war, at times it will deliberately do so. What Ukraine has done, starting with preparing Molotov cocktails to now getting pretty long-range capabilities that can do real damage to Russian assets, is a form of escalation but is also its attempt to stay in the war and win it. However, there is this element of big decisions being taken perhaps on the basis of very bad information, perhaps in some cases deliberately inserted, perhaps because things are a muddle; they are confused. In fast-moving situations, knowing exactly what is going on is very hard. We are much better at that than we used to be. It is important to stress that, by and large, the new technologies help a lot in situational awareness. Commanders now are far better informed than they were. Often, they were completely in the dark and much more vulnerable to rumour and muddle before.

So we should acknowledge where things have become easier as well as where they may get worse. As with all these systems, when they are working fine it is much better. When they make a mistake precisely because we trust them to work well, the consequences can be larger.

Q75 The Lord Bishop of Coventry: Thank you, Professor. That is very helpful, and in many ways reassuring. My supplementary may sound a little left field, but I think it is related.

A number of warning bells have been rung by people in the industry. The latest, and maybe the most credible, was from Geoffrey Hinton last week. There was also an open letter from a number of scientists that was supported by Elon Musk, and now there is talk of a pause. I think we have all felt that that is very unlikely to happen. Nevertheless, is there something that we need to take heed of? If I could make that a bit more specific, my sense is that you have talked about AI being applied in a fairly narrow sense. It seems to me that their warning is that a general intelligence is not as far off as we might have thought. Is there anything to heed there? Is there sufficient anticipation of that eventuality—or probability, it would now seem?

Professor Sir Lawrence Freedman: Obviously this is a big issue, because there has been a sudden moment of awareness, which I think ChatGPT is responsible for, where all of a sudden an amazing capability is accessible to everybody and it is a shock to the system, from writing judgments in court to student essays, maybe even to writing sermons.

The Lord Bishop of Coventry: I have banned it.

The Chair: You need not respond to that, Bishop.

Professor Sir Lawrence Freedman: Sorry, I could not resist. You have this tool, but anybody who has tried to use this particular tool knows that it tends to—I think the phrase is—hallucinate. After a while, when it cannot make the obvious points and connections, it just makes stuff up because it has formed a tentative connection. This is where it can start to be dangerous, because it can create a confidence in what you are getting.

These issues are all part of the same problem. It is very difficult to say that you just pause this, but it is not so difficult to say that there are certain principles and concerns that are very serious, that must be taken seriously and that should be applied and discussed where possible. Obviously some applications can be developed quite cheaply and easily, but really big models like those that produced ChatGPT and the equivalent can be done only by companies and states with very large budgets, so some control can be exercised that way. They need to think it through.

We also need to teach people how to use it. There was a discussion at my college yesterday that was also relevant to the staff college. As I am sure you can imagine, undergraduates get very tempted to get somebody else to write their essays for them. Rather than saying, “You must never touch ChatGPT”, we say, “All right, do that as your first run. Now let’s go through it and work out how reliable it is, where it has gone wrong, what it missed and so on, and use that as a learning aid in itself”. Part of the challenge is to develop a critical understanding of the technology, to make people literate about what the technology can and cannot do. That is an important part of going forward.

It is clear that a general capacity is being developed here. My line—specific problems, large amounts of data—applies. If you are China and your specific problem is social control, and you have something on everybody, as well as facial recognition and so on, AI may work for you, if that is what you want to do. I hope it is not what we want to do, but it is one way in which it can be used, and that is an ethical and legal question that needs to be addressed. Facial recognition is already quite a hot topic, because you can interrogate a database of lots of faces, and if people are not wearing masks and appear there you can pick them out, except that sometimes you pick out the wrong person.

It is a literacy issue. It is understanding that just because it has come out of the system does not mean to say that it is accurate—going back to the rubbish in, rubbish out problem—and we need to develop a critical

understanding of what it can do. That is a more realistic course than just saying, "Stop", because if you stop, you still have everything we have at the moment.

Again, I am not sufficient a technologist to understand how true all of this is, but these really big systems use an awful lot of computing power and electricity. They are dependent on high-quality microchips and, if they are going to develop further, you need more miniaturisation and the development of content, computers, and so on. It is all linked. We have lived through a period where we have just assumed that everything speeds up amazingly. That may not necessarily be true. There may be a point where it peaks. AI was a big thing in previous decades, but then it subsided because they just could not find a way of taking it much further. Then, 10 or 12 years ago, it suddenly took off again. The sheer capacity constraints may also become a factor.

The Chair: My wife is a rural dean, so I will pass on your advice about using ChatGPT for a first draft with Sunday in mind.

Q76 Lord Fairfax of Cameron: Thank you for your evidence so far. It is very interesting. I have a short supplementary to what you were saying earlier about drones in Ukraine, but maybe I will ask my principal question first. As regards prospective regulation, is there any valid distinction to be made between the development and acquisition of AI weapons systems, and their use?

Professor Sir Lawrence Freedman: Yes. That is a standard issue with nuclear weapons.

Lord Fairfax of Cameron: That was the analogy I was thinking of, yes. Are the international discussions going that way at the moment? That was really what I was thinking.

Professor Sir Lawrence Freedman: I do not know. It is very difficult to tell Governments or companies not to carry on developing, but it is not difficult to pose the legal and ethical questions, and the political questions, to them as they do that. There are choices to be made about what sort of applications you develop and how you develop them. Do you go for large numbers of simple drones, or a few big ones? These decisions are not that different from normal procurement decisions and, as I have tried to emphasise, we should not just assume that everything to do with AI takes us into a new and dangerous area. Some of it is really quite helpful and may make for better decisions and less damaging outcomes.

The basic principle is that as long as there are human choices to be made and you are facilitating human choices, any weaponry that we can see developing can be used for terrible purposes, shameful purposes, illegal purposes, or for defence. That is no reason not to develop it and purchase it, but it is a reason to think hard about how it is being used.

The major concern people have is "The Terminator" problem of it going rogue and suddenly deciding that the people you are supposed to be

defending are actually the enemy, and because you are a robot nothing tells you that this is a bad idea, so off you go. However, these are risks in any conflict. Blue on blue attacks happen because the information is poor; all you know is that you look vulnerable, and you decide to deal with the threat.

I do not think one should assume that the class of problems that are developing here are different in kind to the ones we are quite familiar with from the past and which result from the uncertainties and confusion of war and the difficulties of decision-making. The challenge is using AI to make those decisions better rather than worse. I do not think we will get to a situation where people will allow the systems to take offensive decisions on their own. I have indicated that they already take defensive decisions now, but taking offensive decisions on their own will be much harder. You will see swarms of drones being set against enemy systems that are communicating with each other in some way. That seems perfectly plausible to me. It is a new way of doing what armed forces have always done, but I do not think it poses a particularly different kind of problem to the one we have always had. It is just a different sort of technology.

Q77 Lord Fairfax of Cameron: We have talked a lot about drones in Ukraine. Are you aware of, or have there been any anecdotes about, loitering drones, kamikaze drones or other drones identifying a target and then themselves—

Professor Sir Lawrence Freedman: No.

Lord Fairfax of Cameron: So we have not crossed that threshold yet. There has always been a human somewhere in the loop saying, "Right, see this", and they make the decision.

Professor Sir Lawrence Freedman: It just does not happen. All the imagery we have is of operators. They do not loiter for very long either, identifying targets either for themselves or for artillery.

Lord Sarfraz: Thank you. I have found this very helpful. What is the ideal in-house capability at the MoD to make the right procurement and partnership decisions? You talked about the civilianisation of technology. Even though the MoD has some great technologists, many will be attracted to the private sector with their big compensation packages and flexible work-from-home policies, bean bags, sleep pods, granola dispensers and so on, so how should we be thinking about the in-house MoD capability?

Professor Sir Lawrence Freedman: That is an interesting question. It is worth looking at GCHQ, which seems to acquire everything digital as it is, so maybe they will end up playing a more important role.

How does GCHQ maintain its excellence when it has exactly this problem? It is partly because it appears to be able to pay a higher rate; otherwise, you will just not get the people. A number of people have pointed out that some of the very important technology jobs in the Civil Service are

going at ludicrous salaries at the moment that would require people to almost halve their pay if they were coming from the private sector. Cheltenham may not have as big a problem. Creating an atmosphere of intellectual excellence, a feeling that you are doing something important, is what is really making a difference.

There is flexibility in people going in and out of the private sector. Quite a lot of people, I understand, come back to GCHQ after making their pile, if you like, or after doing well enough in the private sector and then feeling that they like the special challenges. It is not that different from running a university: if you want to have a good department, which depends on really bright people who understand something, you have to give them the support and the intellectual excitement that can come with working on these topics, and the feeling that it is an important public service.

How much you need to do in-house is one of the big questions to which I do not even claim to know the answer. A lot of very good technology companies in this country are working hard on this, so maybe you need to replicate what they are doing in-house or work out ways to work closely with them, while accepting that this may raise security issues for some of the people involved, or ethical issues such as whether they want to work on military matters.

It is not simple, but for a lot of areas you really do need to look at public-private partnerships rather than expect everything to be in-house. I just do not think we have the capacity to do enough when there is so much good stuff elsewhere. However—again, this goes back to a very old debate about what needs to be done in-house and what can be contracted out—you still need high-quality people to evaluate what you are getting and to know the right questions to ask, even if they are not going to be doing all the development work themselves. It is vital to have people in government with sufficient authority and competence to be able to assess what they are getting properly, to ask the right questions, and to make sure that these legal and ethical, and political, questions are fed in and that you are not just giving contracts to people to do what they would have done anyway—this can happen—and which they find interesting but may not be serving the national interest.

The Chair: I can guarantee that those questions will resurface when we have Ministers in front of us later in our inquiry.

Q78 Baroness Doocey: You mentioned earlier that you thought there might be a choice later about Governments buying stuff off the shelf. Can you ever envisage a situation where the mindset of our MoD, which I must tell you I do not think of as an exciting and fascinating place, might be to buy off the shelf rather than want to control it themselves?

Professor Sir Lawrence Freedman: Historically, the MoD has not been great at that. There have been a lot of very long procurement processes which, because they are for very specific military requirements, take longer, and the turnover of civil servants and so on means that they take longer still; we know the problems. Buying regularly off the shelf, in

areas where it is appropriate, can speed things up a lot. Again, it is not a new debate. The MoD is capable of buying off the shelf. The challenge is always being very clear about your specifications and what you want, and then letting the contractor provide them rather than keep on changing all that. I do not think it is impossible for them to do. Again, we are talking largely about software here, which is very different from the big-vehicle projects, which are the ones that tend to go on a long time.

Baroness Doocey: Do you believe that if the Ministry of Defence is getting over the problem of paying much more than per its normal scales by using consultants, using consultants who are obviously not members of staff and who may not have the same ideas about loyalty and all the rest of it can mean that the MoD can compete with private businesses that seemingly can pay whatever they like?

Professor Sir Lawrence Freedman: I do not have enough practical experience to give you a very authoritative answer. Unfortunately, consultants are a way of getting around some of the limitations of the Civil Service. Often, they can provide a more expensive service but not always a better one, and in areas like this, which are so important and potentially demanding and exciting, and so on, you want to create the élan that you find at GCHQ. It is not purely transactional.

Q79 **Baroness Hodgson of Abinger:** I want to focus for a moment on the defence of the UK and what capabilities the UK might need to develop to combat the use of AI in automatic weapons systems by enemy combatants as well as to defend its own systems.

Professor Sir Lawrence Freedman: That is part and parcel of the general development of defence policy. There are the standard questions of the defence of the UK, which come back to the fact that fortunately we are an island, which helps a lot. There are big air defence questions, and AI, as we have indicated, may be helpful there. There are questions about expeditionary forces and the sort of problems they may face. However, as I have indicated, we are some way off from the really effective use of autonomous systems. Again, it is just part and parcel of drone and anti-drone development.

There is also a broader security question about the whole cyber area, where, as I have indicated, there is a constant offence/defence duel going on, much of it directed against criminals rather than other states, which AI will make more of a challenge. You can see opportunities developing in relation to the City and the integrity of the banking system—all these sorts of things—which malign states, and certainly criminal organisations, may wish to exploit.

Then there are the other things suggested before, such as attempts to manipulate the information environment and so on. These will increasingly need attention. On the legal and ethical questions about AI development, and the possibilities of developing plausible social media figures who turn out to be complete fabrications, again, these sorts of issues will become increasingly important.

Q80 **Lord Mitchell:** Would the inherent right of self-defence and Article 51 of the UN charter permit states to deploy totally autonomous weapons systems in self-defence?

Professor Sir Lawrence Freedman: Article 51 is a very useful article, simply because it allows you to defend yourself without getting the rest of the Security Council to support you. I do not think it makes any difference to how you do it. You can go back to the just war theory and international humanitarian law for guidance on what to do, but what you decide to do in self-defence depends on the capabilities you have, the challenges you face, and the ethical questions. You have to use what you can.

The issue is whether autonomous systems will be of any use and whether there is any value in them, or whether we want to avoid them. That is still a big debate. I do not think we should see them as the ultimate weapon that you bring into play only when you are faced with the direst threat. If they are going to develop at all, they will be part and parcel of our national capabilities.

I am not sure whether that answers your question, but I do not see it being a special situation in that regard.

The Chair: Professor Freedman, thank you very much. I am sorry that the configuration of the room today means that much of the dialogue has been long range, but you have taken our inquiry evidence on to a whole new level, and my colleagues and I are very grateful to you.