



Artificial Intelligence in Weapons Systems Committee

Corrected oral evidence: Artificial intelligence in weapons systems

Thursday 22 June 2023

10.05 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 Grocott; Lord Hamilton of Epsom; Lord Houghton of Richmond; Lord Mitchell; Lord Triesman.

Evidence Session No. 11

Heard in Public

Questions 139 – 155

Witnesses

I: General Sir Chris Deverell, Former Commander, Joint Forces Command; Tsvetelina Van Benthem, DPhil Candidate in Public International Law, University of Oxford.

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 witnesses

General Sir Chris Deverell and Tsvetelina Van Benthem.

Q139 **The Chair:** Good morning to Tsvetelina Van Benthem and General Sir Chris Deverell. It is a great pleasure to have you with us this morning. From glances exchanged, I think some of the members of the committee may be familiar to you already.

You know the form for this sort of thing, I am sure. There will be a transcript and you will be able to correct any factual errors. The session is being broadcast. We are programmed in to go to about 11.30 am, but we will see how we go. You may feel that you have imparted all your wisdom before 11.30 am, but that has not been the pattern so far in our inquiry.

Welcome to you both. Sir Chris, we had a quick word about this beforehand. Tell us about your role as a mentor for the Creative Destruction Lab at the Saïd Business School at Oxford.

General Sir Chris Deverell: Yes, with great pleasure. The Creative Destruction Lab is a global collaboration between various business schools around the world. It is a not-for-profit mentorship programme for start-ups. Essentially, it brings together academia, start-ups, MBA students and mentors, and hopes that the process of bringing them together benefits the start-up and, in turn, that the start-up can build something massive for the benefit of humanity.

Q140 **Lord Houghton of Richmond:** Good morning. It is good to see you. Thank you for coming along. The committee has reached a dangerous stage where we have got a false sense of our own competence and knowledge about some of this, so forgive me.

One thing we picked up yesterday from some of the language that has been in much of our paperwork is that you can roughly separate artificial intelligence into what you would call narrow, general or super-intelligence. On super-intelligence, the outcomes of the artificial intelligence are beyond human comprehension. Patently, a weapons system that employed super-intelligence would not be permissible on a battlefield, because the outcomes of what it says, does or activates are beyond our comprehension.

However, there are some strong arguments that, in both narrow and general artificial intelligence, so long as an element of relevant human control is retained somewhere in the system of an autonomous weapon, these will be permissible. At the narrow end, if such a thing is set up purely for the self-defence of a ship in quite an obvious battle scenario, when the ability to align with international humanitarian law is relatively straightforward, you could have almost full autonomy. In the general area, there is a question about the level of effective human control of the system that should still be retained.

First, do you agree with the analysis that, for the foreseeable future, some form of effective human control will have to be a part of fielded AI systems? Secondly, drawing on your background, might you suggest

what you consider to be the training burden, the command burden or the human competence burden this might place on the people we have in the Armed Forces?

General Sir Chris Deverell: There is a lot in that question, so I will try to parse it in a way that I feel competent to answer. The classical taxonomy in academia regarding AI is between AGI—artificial general intelligence—on the one hand and other forms of artificial intelligence on the other, which are often described as narrow. I am not really familiar with the third category. I get what you mean by the statement, but the term “artificial general intelligence” is generally taken to mean the ability of artificial intelligence to do things better than humans could do, and not just a specific task but any task.

We are quite some time away from that state in the development of this technology. There are debates about how far away we are from AGI. I recently heard an Oxford academic say five to seven years, but just before GPT-3 people would have said 10, 15 or maybe 20 years away. That state of artificial intelligence, where it can do better than humans at any task rather than being bespoke for specific tasks, is not with us yet. If that is what you mean by super-intelligence, the first thing to say is that it is a while away.

We have, however, had various forms of AI that can do specific components of a task. For example, we have had computer vision for a while, which can classify objects that it detects on the basis of a prediction as to the correct label to apply to an object. That is what it is doing: it is a prediction. From a collection of possible labels, what is the appropriate one to apply? If you have labels for battlefield systems, is the correct label for the thing I see “tank” or “bush”?

That is the first thing. I am not certain I understand what you mean by these different definitions. Secondly, we are a while away from an artificial intelligence that can do any task better than a human.

As to human control, if you want to be compliant with international humanitarian law today, you absolutely must have human control. We cannot yet have sufficient confidence in these systems to say that they do not have the possibility of mistakes and error. What you are controlling for with that human control is the possibility that the machine might make a mistake, in my opinion, rather than the notion that the machine should never be allowed to make this decision or prediction.

I can conceive of a world in which—you can make this true by definition—a weapons system has 100% probability of correctly identifying and locating a target, 100% probability of selecting the appropriate weapons system to use against that target and 100% probability of then hitting that target in a way that is discriminate, proportional, necessary and so on—in other words, compliant with international humanitarian law. That is a conceivable state. What we are trying to deal with is a situation where we do not have 100% probability of those things today. Therefore,

it is essential to have some human control in the decision to use lethal force.

You then asked about training. We currently train our people to think about the characteristics of the weapons systems they are using and to make decisions based on those characteristics and on their knowledge of international humanitarian law, the rules of engagement they are under and so on. Fundamentally, that is not a new task.

I believe it is possible to codify what the weapons system can and cannot do in such a way as to make it explicable to a commander or a user on the ground and for them to then employ those weapons systems with a very low, albeit not zero, probability of error. It would be a new weapons system and we always need to train people in new things, but I do not believe it is impossible or unduly excessive. Does that answer your question?

Lord Houghton of Richmond: It does. Forgive me about the definitions, but definitions are an elusive and changing thing throughout this. There is no great commonality about definitions.

Tsvetelina Van Benthem: Your question is positing problems about the empirical and technological state of affairs, but also legal questions about permissibility and questions of control under existing law.

I would like to start with a little note on the use of terms. The terms "artificial intelligence", "autonomy" or even "human control", for example, are not terms that feature specifically in the law of armed conflict. What we have are specific obligations that require different types of decision-making, which entail different forms of judgment in relation to the way we choose our targets of attack and the way we choose the modalities of reaching those targets.

In that sense, the discussion about human control places us in a position where we can agree on the need for it but fundamentally disagree on what we mean by it. We have observed this in the intergovernmental discussions on autonomous weapons systems. There is an agreement about the centrality of the human element. In the recent joint statement by 70 states that was made last year to the UN General Assembly, they spoke about exerting appropriate judgment, control and involvement, but there was an acknowledgement that different states will have very different understandings of what is appropriate when using autonomous weapons systems.

In this sense, to answer your question, we would have to go to specific obligations under international humanitarian law. This is a better approach than seeking some kind of overarching requirement of human control that may not exist under humanitarian law. This requires a more rigorous interpretation and analysis of treaties, and identification and analysis of custom.

The key question that was already addressed by Sir Chris relates to factual uncertainty coupled with legal uncertainty. How do we deal with fog of war issues when we are not entirely certain about the content of the obligations under the law of armed conflict? Do they require intention? Can they be breached by recklessness or negligence? The scope of unpredictability and unintentionality is the sharp end issue in this discussion.

- Q141 **Lord Triesman:** Thanks to both of our witnesses. I have a couple of questions to try to peel back some layers of this onion. Apart from the development of different kinds of AI, a number of witnesses have made it clear that on a contemporary battlefield things happen at much greater speed. There is a need to be able to respond to fast events in part because they happen so quickly and, when we are looking at operational environments, we need to take into account the fact that an enemy would be attempting to deceive us as we would probably be attempting to deceive them as well. A relatively short time into an engagement, events will be happening quickly and there will probably be attempts at deception.

This may come back to that question about the specific character of the engagement. I can get that point, but in what ways is it possible to comply with international humanitarian law when the variables are changing so rapidly? What is the basis on which we should feel we can do so? I probably ought to add that I am not saying we should deny ourselves the capability of either defending ourselves or, if necessary, attacking other people, but how do we achieve compliance with international law in these circumstances?

Tsvetelina Van Benthem: The first point in relation to this question is that the discussion on autonomy is exacerbating these concerns, but this dynamic is already present in any other type of military operation. Take, for example, the difference between deliberate targeting and dynamic targeting. We already have measures that make it more difficult to comply with these obligations and ensure full protection of the civilian population, such as when there are attempts at deception or when combatants are intermingled with the civilian population. In many ways, autonomous weapons just expose the concerns we have more generally in relation to military operations.

When it comes to specific obligations, there is no difference in the requirements of the obligations, whether we are talking about environments that are fairly controlled and not unfolding at a fast speed or environments that are dynamic, but we have to apply these obligations in a contextual manner. For example, when thinking about the presumption of civilian status, which is binding on states that are parties to additional protocol 1 of that treaty and which already exists under custom, the way we form doubt about status will inevitably be different depending on the position we are in and the context in which we are applying this rule.

In that sense, to your question, we still have to comply with these obligations, but the context will determine what exactly is required by parties to conflict in the specific circumstances of use.

General Sir Chris Deverell: I totally agree with that. I would add two points before answering the question, if I may. It is often the case that we compare a new technological solution with perfection rather than with the current reality. The fog of war exists today in very large measure. We have to deal with it and in a way that is compliant with international humanitarian law. There is at least a possibility—some would say a likelihood—that the use of artificial intelligence will reduce the fog of war to a point where its ability to discriminate, for example, is better than a human's.

To answer Lord Triesman's question, what would it take to do that in an evolving operational environment? It would require a very thorough understanding of the capabilities of the weapons system; an ability to update the weapons system to take into account new circumstances, something that is not currently available with any weapons system we have; and an ability to decide whether to use them or not, to make a conscious choice, if the operational environment had changed in such a way that would generate a risk that their use was not compliant.

None of that is impossible in principle. One of the great benefits of AI is that it can, for example, be trained and retrained to reflect new data, provided one designs a means to embody those changes in frequent software updates on the system. As we discussed a bit earlier, educating the military on the factors that would affect a decision to use or not to use a particular weapons system is also not impossible. Indeed, it is something we do already with existing weapons systems.

Of course, the fact that I believe these things are possible does not make them easy. If you wanted to have an AWS, which is not currently UK policy, we would have to do a lot of work in modelling and simulation during the design of the system, extensive trialling during development as well as a lot of user training. We would have to monitor in-use and have regular reviews. There are many components of a decision to have such a weapons system, but, to answer the original question, I do not believe it is impossible to adapt the parameters the weapons system is using to reflect changing operational circumstances. Indeed, it is possible that a machine will be better at that than we currently are today, absent such technology.

Lord Triesman: We have been talking about this in the context of the fog of war. Are there similar implications during the fog of peace, when people are perhaps making cyberattacks or closing down fundamental systems using other kinds of technology?

Tsvetelina Van Benthem: Yes, absolutely. This has been discussed in other fora in relation to the regulation of information communications technologies. There are questions about how we think of attribution, for example, where methods of deception have been used. That even has

consequences when we think about enforcement because states may seek to adopt countermeasures towards something they see as a violation of international law, but they may be mistaken because what they perceive as coming from one state may not be coming from that state at all, but from a non-state actor.

As you say, this is not a question that relates specifically to conflict. It is a more general question about how technology can enable the masking of identity in certain ways, which makes it more difficult for us to see the facts behind a particular operation.

General Sir Chris Deverell: Yes, I agree. Furthermore, there is no hard dividing line between peace and war; there is a continuum. At levels below major conflict, people may use technologies against us or we against them that need to be governed by the law. Traditionally with autonomous weapons systems we would talk about systems that employ lethal force. In the main, cybertechnology is not directly doing that. Nevertheless, to the degree that it could, it would be governed by the same principles.

Q142 **Lord Browne of Ladyton:** I will try to ask a question that makes this more understandable. I am not entirely sure that I will succeed, so forgive me if I do not.

From my perspective, part of our responsibility in doing this is to find a way of explaining the evidence we receive in a way that allows executive politicians to get the trust of our publics to use these weapons systems. This is an incredibly difficult challenge. The danger is that it will be avoided because it is so difficult and that we will just move on. We already know from polling that there are very few publics, if any, who trust AI systems at all, never mind AI weapons systems.

Let me just ask a combined question that comes from things both of you have said. You left us, Chris, with this statement that some kind of human control is presently necessary. We probably all agree with that, although we probably do not really understand why we do. That range of human control, although not infinite, is very diverse and wide. I consistently hear people say, "It will always be under the control of a human because you can switch it off". That is not what we are talking about.

Tsvetelina, you left me with the impression—I want to know whether this is what you are trying to convey to me—that, until we are in a position to codify with some significant detail what these contexts will be, we cannot be compliant with international humanitarian law. Therefore, we should just park this until we have done all this detailed work.

I ask you this question against this reality. In every aspect of law—I did not practise international humanitarian law, I have to say—there are concepts that are protean; they change. Five years later, a judge will say, "This is what breach of the peace means in the current context because of the way society is developing". We have not needed to anticipate all these changes.

If you are educated in a Roman law-informed jurisdiction, probably common law, you end up doing what we do in Parliament. We try to legislate for every possibility, which means that we have rooms full of legislation that is impossible to apply. Everybody thinks that unless their specific issue is being dealt with, the issue is being dealt with only in words, if at all.

Would it not be possible for this concept to develop through a judge-led interpretation of what people like us put in legislation rather trying to anticipate everything in advance and not using these systems until you can? Have I misunderstood you? Am I near what you meant?

Tsvetelina Van Benthem: As you spoke, I was thinking that there are probably 11 points I will have to make. I will try to combine them in a more meaningful way.

You underline a dynamic that is common to all law. Maybe we cannot, and maybe we should not, be prescriptive about all the circumstances to which the law will apply. At the same time, we cannot deal with areas where there is significant uncertainty about the law. That is what we see with international humanitarian law. It is not just me saying this. There are states coming forward with this position at the United Nations in their national positions.

There is a concerted effort to clarify the scope of international humanitarian law, which does not need to go to the other extreme of completely specifying all the applications of that law, but we at least have to be clear about the elements of the different rules. For example, can you violate the prohibition on attacking civilians only through intention? Can you be reckless when you violate it? We could speak about the presumption of civilian status. What does "doubt" mean? What kind of certainty or reasonableness standards are embedded in that rule?

Unless we have some reasonable sense of what the law is, we cannot implement this at home. Sir Chris spoke about all the procedures for the training of forces domestically. It will be very difficult to do that in a meaningful way when it comes to these emerging technologies unless we have a better sense of how they apply to autonomous weapons systems.

Lord Browne of Ladyton: For homicide to be murder, it has always involved either mens rea or recklessness. Why should it be any different in international humanitarian law? Why should you distinguish between these two?

Tsvetelina Van Benthem: In international law, states have agreed to different rules that contain different subjective standards. Some of them contain no subjective standards at all. It does not have to reflect domestic law. We are speaking about two different systems.

You are speaking here about individual criminal responsibility under domestic law. When we speak about international humanitarian law, these are obligations imposed on parties to conflict. In fact, for many of

the obligations we have in international law we are not required to have a standard of intent or recklessness; we may have strict liability or negligence. In that sense, we should not be thinking about the convergence of these regimes. There are separate regimes and separate rules, and we have to do a rule-by-rule analysis.

To your question—you pointed out the questions of methodology before—it will be really important to have a strict methodology of what exactly we are talking about, which legal regime and which elements of that legal regime? What is the modality of clarification? This will be the key point, as we have seen in the debates. Some states, such as the United Kingdom, insist on a clarification exercise for existing law. Other states want to move forward with a new binding instrument that will specify existing law and maybe build on it. This is a second and subsidiary question to the one you are asking, essentially, about the goal that is in front of us, which is some form of clarification.

General Sir Chris Deverell: Trying to codify every scenario and say, “This is the right thing and this is the wrong thing to do in that scenario” is just unworkable. At some level we have to rely on principles. To the degree that we can agree among states on clarifications of the existing principles, that is so much the better, because they can then be more easily explained than an abstract principle.

Coming back to your original point, Lord Browne, you talked about people not trusting these systems. That is definitely true, but it depends on what you mean by “these systems”. At some level, a political decision describes the standard we will accept for risk in the field of driverless cars, for example. It is perfectly possible to imagine a world in which driverless cars are much safer than cars driven by humans, but, rightly, we hold the machine to an extremely high standard.

I would apply the same principle in the use of lethal force. We should have very high standards before we entertain it, but it is not impossible to conclude that AI in weapons systems might end up being beneficial in terms of accuracy, discrimination, proportionality and so on.

Q143 **Lord Hamilton of Epsom:** As I understand it, Tsvetelina, under international humanitarian law, the sanction is war crimes. To get other sanctions that come below war crimes you would have to get agreement from a lot of nations that will not give it. We are talking about a hypothetical that will happen. We are left with international humanitarian law, which says that you can charge somebody with a war crime, but you cannot charge them with anything else. Am I right on that?

Tsvetelina Van Benthem: We have to distinguish between different forms of responsibility. When we speak about war crimes, we speak about the responsibility of individuals under international law for certain serious breaches of international humanitarian law. For example, this regime can be operationalised at the international level, subject to there being a forum that has jurisdiction over these crimes, or the domestic

level, either in the state of nationality or territoriality of the crime or on the basis of universal jurisdiction in a third state.

We also have the concept of state responsibility, which is key for what we are talking about here. This goes back to the question about trust in Governments when adopting these systems. This is separate from the idea of individual criminal responsibility. State responsibility is the responsibility that entails automatically under international law when a state violates its obligations under that law.

In that sense, we have to transition to the general law of state responsibility about how we can implement that responsibility. There are different ways of inducing compliance with the law. Countermeasures are one example. For example, when European Union states impose sanctions on Russia in relation to its aggression in Ukraine and those sanctions breach obligations that these states have vis-à-vis Russia, these are countermeasures.

If we have a state that is breaching international humanitarian law through the use of autonomous weapons, we have different tools in the international system to get that compliance. It might be an invocation of responsibility, subject to there being a court or tribunal with jurisdiction; decentralised enforcement through countermeasures; some institutional framework, such as reporting obligations or review mechanisms; or even naming and shaming, which will have reputational costs.

We have to think of enforcement and implementation of responsibility in a structured way, across the different stakeholders and the different ways in which they can be influenced such that they comply with the law.

Lord Hamilton of Epsom: For this to be classified as international humanitarian law, you would have to get agreement across the board that you will do that. Is that going to be forthcoming?

Tsvetelina Van Benthem: This is already the case. Let us take a clear example. If a state were to use an autonomous weapons system to intentionally target civilians, that state would already have violated international law. In that sense, we transition to the different enforcement mechanisms that I have mentioned. We do not need an additional agreement.

The agreement is already there. The obligations are the existing obligations under humanitarian law, human rights law and customary international law more generally. We also have other ways to implement responsibility that we can use to induce compliance with those obligations. For responsibility to occur, there is no separate step for which we need additional consent.

General Sir Chris Deverell: Tsvetelina quoted the example of a state using an autonomous weapons system deliberately to target civilians, which would be illegal. It is not the fact that it is an autonomous weapons

system that makes it illegal; it is because targeting civilians is against international humanitarian law.

I completely agree with the rest of what Tsvetelina said with the exception of one word, which I think is important in answering your question. Tsvetelina said that there are means to get compliance. I think there are means to seek compliance. These tools can be very effective, but they are not a way of totally ensuring the outcome you seek. You can impose sanctions, for sure. You can pursue people in courts. You cannot be certain that these mechanisms will have the outcome you seek, which is for people to be held responsible for their actions.

Lord Hamilton of Epsom: It has to be said that the drone pilot who took out a lot of people at a wedding party in Afghanistan never faced any charges at all.

General Sir Chris Deverell: I do not know the circumstances. It would be important to look at the circumstances.

Tsvetelina Van Benthem: There is a difference between there being a violation of the law and there being an attempt to enforce responsibility for that violation. Very often the latter will not happen because there is no good forum and no one who has the interest to take the case. That does not mean the violation did not occur.

I am not commenting on the specific case, but there are many instances where presumably the law has been violated and there is evidence that the law has been violated. Yet, due to a variety of procedural limitations in international law, no responsibility has been implemented in practice.

Q144 **Lord Clement-Jones:** This brings us back in a way to the definitions. We talked earlier about the definition of AI, narrow AI, AGI and so on. When you come to autonomous weapons systems—AWS—there is no single internationally agreed definition. The MoD does not have an operative definition of AWS. Historically it did, but it included the word “understanding”, which always struck me as being slightly strange.

How do you define AWS? Is it important to have a common definition and, if so, why? Should we adopt a working definition of AWS very soon going forward?

General Sir Chris Deverell: I think of AWS as, first and foremost, a combination of hardware and software designed to independently detect, analyse and engage targets, including with the use of lethal force, without direct human control over the individual targeting decisions, although humans will of course have been involved in setting the parameters and thresholds within which the system operates, during its design.

Lord Clement-Jones: That is very similar to the UK joint position submitted to the GGE.

General Sir Chris Deverell: It is a slight expansion on it, yes. To answer the second part of your question about whether we should have a working definition, there is something slightly odd about having a policy not to have something without defining what it is. Furthermore, there is another really important reason for having a definition of AWS, which is not necessarily one that people have drawn your attention to so far. If you could have a definition of what is an AWS, it would give you a definition of what is not.

It would also stop you slipping by accident, which happens a lot in discourse and, indeed, in some of the evidence in this committee, from a world in which you are saying, "We don't want an AWS" to a world in which you are saying, "We don't want AI". The latter would be like saying, "We don't want to use electricity because you can hurt people with it".

Lord Clement-Jones: Could you speculate on why we have not adopted a working definition to date?

General Sir Chris Deverell: There are some good and some less good reasons. One good reason is the feeling that, if we leap too early to a concrete definition, we will discover in the evolution of the technology something we did not anticipate, and therefore we should wait a while before we reach a definition so we do not by mistake either include or exclude something that we ought not to have. That is one good reason.

There are other perhaps less good reasons. This is arguable, but I suspect the UK would ideally like to come up with a definition that had general international agreement. It does not want to be in conflict or in contrast to, for example, the United States on this, if it can avoid it. It has not yet proved possible for the key like-minded allies of the United Kingdom to come to an agreement. That is partly what it is.

Lord Clement-Jones: It is not about keeping our options open as regards the impact of international law.

General Sir Chris Deverell: There is discussion going on—Tsvetelina referred to it—about what all this means and does not mean. It is two sides of the same coin. "Keeping our options open" is a slightly pejorative way of describing it. "Being thoughtful" is an alternative way of describing it. We are being careful not to do something by accident.

Our understanding of this technology is evolving very quickly, but it is time for the UK to have at least a working definition. That would not necessarily state, "This is what it means for ever and a day", but rather, "When we say we have a policy not to have automated weapons systems, this is what we mean by an automated weapons system".

Tsvetelina Van Benthem: I agree. Going back to your reference to the MoD's submission, although there is a convergence among state actors as well as international organisations about at least some characteristics of these systems, they will use slight variations in the wording. For

example, in the submission the precise wording is “without further human intervention subsequent to deployment”. Even though this definition answers some questions, it leaves quite a few as well.

For example, the way states think about the content of “further human intervention” will differ. To some of them, simply nominal human input, perhaps on the basis of automation bias, would not count as human intervention. Thus, even if there is a possibility to adopt this type of intervention, it would not mean it is not an autonomous weapon. It would still qualify as an autonomous weapon because it is simply nominal human input.

That said, in terms of desirability, I can see the pitfalls of adopting a definition. Lines that make sense today may make very little sense tomorrow. There is a risk of arbitrariness when adopting a definition, but at the same time the characteristics we have converged on give us a good enough sense of the types of risks we worry about, which are about some type of further distancing between battlefield outcomes and direct human intent. It is this idea of intermediation: how the intent is being intermediated from the decision by the human up to the battlefield. In that sense, it seems to me that the definition provided by Sir Chris, and the one that states converged to, is a meaningful one.

Lord Clement-Jones: Despite the fact you are just a bit concerned about going nap on a definition now, as a working definition, you would be happy with that and you think it is desirable that we have something like that going forward in order to get our heads around some of the legal liabilities and so on.

Tsvetelina Van Benthem: Yes, we can phrase it as a common understanding of characteristics. It is not necessarily a definition in the way we can define a chemical weapon or a cluster munition, but it is still some form of common understanding on characteristics and elements.

Lord Clement-Jones: Could you come up with a better definition at this stage?

The Chair: You can write to us afterwards, if that helps.

Tsvetelina Van Benthem: The one element that slightly worries me about this definition or common understanding is the notion of intervention.

Lord Clement-Jones: Yes, the extent of human intervention.

Tsvetelina Van Benthem: In that sense, with a little more time to reflect properly on the wording, I will be happy to come back. This is the element we need to work on to have a better sense of what we are dealing with.

Q145 **Lord Fairfax of Cameron:** I have a supplementary question, and I will try to make it as clear as possible. If I fail, I apologise in advance. Against an extremely fast-moving tech background and in the absence of

international agreement—we all know these negotiations have been going on for a very long time—I want to distinguish between compliance with international humanitarian law and the involvement of meaningful human control.

Consider a situation where the outcome, under the current situation, was in compliance with international humanitarian law, but the party in question had not exercised any meaningful human control. Some states—we can all think of who they might be—are thinking about investing a whole lot of money in this at the moment. In that scenario, you could have weapons systems that are almost completely autonomous.

In its written evidence, the MoD said that it wishes to maintain, as far as possible, a competitive advantage. Would we be tying our hands by saying that we will always have meaningful human control, even if the outcome is in compliance with international law? Are we going to maintain that position? I am sorry if that was not completely clear, but I think you understand some of the points I was trying to make. I just wondered whether you had any comments.

General Sir Chris Deverell: You make a key distinction. Tsvetelina has already made this point. Human control is not explicitly required to be compliant with international humanitarian law.

We have no choice about the latter. As a state, we have adopted and strongly recommend that others pursue compliance with IHL. To the degree that we deem that a weapons system could not be used in a way that is compliant, we cannot use it. It is not a choice whereby we ask, "Should we have this AWS or not?", because it gives us a tactical or even a strategic advantage. If it is not going to be compliant, we cannot use it.

Lord Fairfax of Cameron: You mean if it is not compliant with our ethical system, not international law.

General Sir Chris Deverell: If it is not compliant with our own ethical systems and with international humanitarian law, we cannot use it. I will elaborate on your point. The absence of meaningful human control does not necessarily make it non-compliant.

Lord Fairfax of Cameron: Yes, absolutely, that is my point.

General Sir Chris Deverell: It may be a distasteful thing, certainly in some quarters of society, to imagine a world in which machines can kill people, so we will have to think about it very deeply, but in and of itself it is not automatically non-compliant. You can envisage a world in which it could be compliant and we could choose to have them, such that our policy of not having them would put us at a disadvantage. The issue then becomes about how big that disadvantage is. To me, you are talking about a fairly extreme use case here. We are quite a long way from strategically telling numbers of autonomous weapons systems.

Lord Fairfax of Cameron: What about mass drone strikes or a thousand little hornets?

General Sir Chris Deverell: Compared to the destructive power of nuclear weapons, for example, this is just of a different order. I am not saying that that world cannot exist at some point in the future, but right now, today, we are quite a way from it.

Lord Fairfax of Cameron: One reason why I asked is because we have been very conscious in our deliberations that our report must not be out of date on the day it is published. This is such a fast-moving area. I was trying to look forward, drawing on your great experience, and anticipate whether we might one day have to change our minds. What are your thoughts? There are no easy answers.

Tsvetelina Van Benthem: Yes, it is a fantastic question because it asks us to think about the degree of convergence between the content of the law and this idea of meaningful human control. My issue with this is that the question almost assumes that we operate on a common understanding of what "meaningful human control" is, and I am not sure that is the case.

For example, one can envisage "meaningful human control" as the programming stage where you specify the parameters of action and you train your personnel in how to use it. That is it; that is meaningful human control. Another could think of human control as the requirement of explainability throughout use, as geographical or temporal confines or even as a necessity to have direct human input before every individual use of force. If that latter end of the spectrum is what we understand by "meaningful human control", the standards will indeed diverge. We will have something that is required under international humanitarian law, and then the idea of meaningful human control will require something additional.

One thing I wonder about is the extent to which that discussion about meaningful human control and its specification will feed into the way we interpret the rules of IHL. Will the two gradually become one? Will it have the opposite effect, in that states will consider that this is not a workable definition or concept and diverge from it? As of today, because we do not have a clear understanding of what meaningful human control is, it is very difficult to answer that question.

Q146 **Lord Hamilton of Epsom:** We have already dealt to some degree with this question, which is about overcoming the problems with changing international humanitarian law. This question is for Tsvetelina and is about the Convention on Certain Conventional Weapons, which I am sure you have seen. Before I ask it, I would like to give you a line out of the briefing we were given against this question: "Decisions at the CCW are taken by consensus, resulting in deadlock due to opposition to a ban from several countries, including the US, Russia, India and Israel". Do you agree? If you do, it makes the question pretty otiose.

Tsvetelina Van Benthem: "Deadlock" is a very strong word. Even at the May session this year, where there was a sense of disappointment in the room, a consensus document still emerged in the end, which was some

form of agreement on autonomous weapons as they relate to international humanitarian law. It is true that divergent opinions on what the law is lead to impediments in reaching common understandings on the law, but I would not go as far as saying that there is a deadlock.

I suppose I see where the question will go, which is about further steps in the CCW and the possibility of using the CCW as a viable avenue for discussions. This is still the case. Despite geopolitical differences—as you can imagine, since February of last year it has become much more difficult, not just in Geneva but in New York, to reach agreement on issues—and on the basis of the different capacities of states, there is still a lot of incentive across all stakeholders to see where the CCW discussion can go.

Lord Hamilton of Epsom: Let us hypothesise that we do not get international agreement; I personally think it is unlikely. Should we be then looking for some national solution to this? That is something we could achieve. What are the risks that we would put ourselves at a disadvantage if we ever went to war against enemies who would not be abiding by our ethics and morality on this?

Tsvetelina Van Benthem: Just to clarify, does “agreement in the CCW” mean a new legally binding instrument or an agreement on the clarification of the law? In some ways, that is what we have had in the past years at the GGE. Is your question going towards the possibility of an international legally binding agreement?

Lord Hamilton of Epsom: A lot of effort has been put into getting this, and a lot of the evidence we have had from other witnesses says that we are very unlikely to reach any conclusion, because at the end of the day the main players will not go along with anything that is too restrictive. I do not think we will get international agreement, but it is more plausible to say that we might get national agreement as to how we deploy AI and automated weapons systems.

The question is whether, if we impose those restrictions on ourselves, we would be putting ourselves at a disadvantage with other nations that are not inhibited by these restrictions.

Tsvetelina Van Benthem: In principle, yes, to the extent that we are thinking about national restrictions that would go beyond what international law already provides. We can think of national measures that are just meant to implement the obligations that we have. In that sense, the fact we have national measures about this and another state does not have national measures does not mean this state would not be breaching international law if it were acting in a certain way.

That said, there may be a little bit of an exaggeration about the degree to which these national measures will place us at a disadvantage. In fact, states have an interest in deploying predictable and reliable systems. This is not just about the risk to someone else’s civilian population; this is about the risk to their own troops. This is about the soldiers who are

deployed in battlefield contexts and the different systems that will operate around them.

It is also about building public trust, which is a question that came up before, and national measures can enhance that public trust. Even if the national measures go beyond the law, they will not necessarily place us at a disadvantage. In fact, they may generate much more trust and better procedures to achieve exactly what we want through the use of these systems.

Q147 The Lord Bishop of Coventry: Thank you both for your contributions today. Tsvetelina, may I also thank you for your written evidence? I found it very helpful, and I am sure others did as well.

You may feel you have said enough about the first part of the question, but you sound a little more hopeful than perhaps some around the table about the negotiations that are happening through CCW. I wonder whether there is anything else you might want to say about your observation of the work of the Group of Governmental Experts and what sort of confidence you have in it.

The second part of the question follows on from that, but you might want to deal with it separately. Very helpfully, you have both been talking about clarifications and, to a certain degree, operational implementation. What does this look like in context? I have had a quick look at the Oslo manual on select topics of the law of armed conflict, the chapter on remote and autonomous weapons. I found that really helpful because I could see its application to particular weapons systems. Is that approach helpful? How would that relate to law? What status would it have?

Tsvetelina Van Benthem: Those are three big questions. To the first one about me being hopeful, this relates to the reason why I wanted that clarification on the previous question.

I am not hopeful that the CCW will necessarily lead to a new legally binding protocol. I am more hopeful that there may be more agreement through consensus reports on the content of international law as it already exists. There is still momentum by states to use this as a convenient platform for dialogue on questions of international humanitarian law. This is why I have some hope, not necessarily for the outcome of a new legally binding instrument, but for continued conversation in that forum.

As for the dynamics of the group, you can see from the reports from 2019 up until now that there seems to be dwindling consensus on substantive issues. Part of this may be because it was easier to reach some type of consensus on the minimum common denominator and now we are struggling to reach something additional. It may also be because there is now a very different geopolitical climate, as we discussed previously.

One of the questions about what will happen relates to what states want. There is a large group of states that want to have a new legally binding

instrument. They are not confined to the CCW. They may decide to take this outside of the CCW. In fact, every disappointment in the CCW may be adding a little more to the incentive to do that. This has already happened, for example when cluster munitions were taken outside the discussions of the CCW, and the Oslo process led to the adoption of the Convention on Cluster Munitions. You would have freedom there in terms of procedure.

It may not be consensus. You can have majority voting; you may have a deadline for the conclusion of negotiations. My worry is this. First, if we reach the point of certain states looking at the legally binding instrument, this will exclude key players. Secondly, what message does it send about the content of international law as it currently stands? Would that in some way erode what we already have by suggesting that what we have is insufficient? There will be an important question about the way we communicate about the efforts that we have.

On your second question of clarifications and operational implementation in the Oslo manual, I agree that these types of manuals are really helpful. We also have the Tallinn manual that has been developed in relation to cyberspace. One part of the question is to ask, "What is the role of non-state actors in clarifying the law, and what will be the buy-in from state actors by reference to these initiatives?" This is a broader question about inclusivity, which in many ways you are showing with us today, in having inputs from other actors that are thinking hard about these questions. The UK has actually called for a manual on international humanitarian law that would clarify its application of autonomous weapons systems. This is a fantastic idea.

Q148 The Lord Bishop of Coventry: May I pick up another question on something you said earlier, that there is dissatisfaction and frustration with the CCW? You hinted that other states were saying that perhaps the whole thing ought to come out of that forum into other international forums. Do you think there would be value in other forums acting as alternatives for some form of regulation? If so, how might the UK play its part in those?

Tsvetelina Van Benthem: First, we have to define our goals with this. What exactly do you want to achieve through that process? I worry that we are not entirely clear on the goals and which modality can actually reach the goals that we are seeking. My concerns are very much those that I mentioned earlier. Many key players will decide to stay out, because the push is to have much more specific regulation in ways that do not comport with the interpretations of existing law that many states already have. It is unlikely that they will want to sign up to this instrument.

Now, would an instrument of this kind then exert some kind of gravitational pull, in that over time states will decide to become parties, or will start complying with its provisions even without being parties because they think that is actually a good way to regulate? With autonomous weapons I doubt that this will be the case. Because there is

such heavy investment in these weapons, my sense is that there will be an aversion, or at least a cautious approach, to very prescriptive regulation.

We focus much of this discussion on humanitarian law, but there are other forums that are currently looking at questions of autonomous weapons, such as the Human Rights Council. The framework of international human rights law, which also applies in times of armed conflict in a complementary way, has to be explored a little further.

The Lord Bishop of Coventry: Do you think this committee ought to give attention to those other obligations under international law, including human rights?

Tsvetelina Van Benthem: I do.

General Sir Chris Deverell: An initial part of the question was to ask, "In the absence of international agreement, should we do something in the UK?" There you have to distinguish between law and policy. We have to have a policy. The absence of a policy would be a policy, so we need to consciously decide, in the UK, how we want to behave in relation to these weapons systems.

Whether you need legislation is a different question. Policy is clearly easier to change than legislation. Legislation appears to be more firmly footed and may have a more positive influence on public trust than just policy. It becomes a question of decisions about trust as to whether UK legislation is necessary in addition to policy.

I do not see the likelihood of states agreeing to something in another forum that they will not agree in the CCW. I do not understand how that would arise. There may be subtle differences, but fundamentally there are states that do not want to prohibit the use of autonomous weapons systems, and I do not think they will sign up to treaties that would force them to.

I hear people making, as indeed some of the witnesses to this inquiry have made, a statement that is in conflict with itself: "These things cannot possibly be developed in a way that is compliant with international humanitarian law, but we need new instruments of international humanitarian law". I do not see how you can reconcile those two statements. I believe that current law can be sufficient to govern the use of these technologies, and I do not believe that a new instrument is required. Clarification is definitely valuable if it can be achieved, but I do not believe that a new treaty is necessary.

Q149 **The Chair:** Can I follow up on the Lord Bishop's question on CCW? We are all familiar with these sorts of international forum. A lot of them are very depressing and there is a great deal of international posturing. States want to be seen for what they are saying rather than what they are doing, but there is a possibility in an environment such as that of creating bilateral or trilateral agreement, which you can then import into

the plenary. Is there any sign of that happening in the CCW or equivalent fora? Who are our friends in what we are trying to achieve?

Tsvetelina Van Benthem: There is not, to my knowledge, in binding legal agreements, but in the submissions from working papers you can already see certain alliances that have emerged. For example, the paper referenced previously was by the UK, the US, Canada, Japan, the Republic of Korea and Australia. In specifying what they considered the law to already provide for, they adopted this two-tiered approach, which we have not yet mentioned in this session, that some types of autonomous weapons systems cannot comply with international humanitarian law and are therefore prohibited. Those that are not prohibited should be regulated. We can see that there is some form of emerging consensus in blocs, but I have not seen this necessarily leading to some kind of bilateral or trilateral binding commitment.

The Chair: Sir Chris, do you see any cause for optimism in this area?

General Sir Chris Deverell: I have nothing to add to what has been said. I do not fundamentally believe that we will get a new treaty, but I do not think that means that progress is not possible in this area.

Q150 **Lord Mitchell:** That sort of kills my question: no treaty, no enforcement. But we must look at enforcement. How could any limitation or prohibition on AWS or the use of AI in defence more broadly or practically be enforced, were we to get there?

General Sir Chris Deverell: It depends on whether you are talking about UK legislation or some kind of international treaty.

Lord Mitchell: It has to be international, does it not?

General Sir Chris Deverell: No, not necessarily. We could legislate on our own say-so, so to speak, and decide to make decisions just for the UK. They would be enforceable in the context of UK law in the courts, but what you are driving at is international agreement. We all agree that the likelihood of a new treaty is low, but we are also saying—certainly from this side of the table—that that does not mean to say that there are not ways of states concluding that behaviour is against existing humanitarian law and choosing not to do it, and equally choosing to pursue other states that do so.

I would not tie the establishment of a new treaty with enforcement. It is perfectly possible to use existing law to pursue states that break it, not necessarily with certainty, as I have said earlier, but as a mechanism for seeking compliance or seeking punishment. That exists today.

Lord Mitchell: Would China see that the same way?

General Sir Chris Deverell: We could choose to pursue China in certain ways if it failed to live up to its obligations under international humanitarian law. Whether they would be successful or have any impact on China is a matter of debate, but it does not stop us concluding that it

has behaved illegally and seeking to take action to try to overturn that or to sanction it for doing so.

Tsvetelina Van Benthem: Your follow-up question makes me very happy, because this is exactly what I was going to come to. Enforcement also depends on what exactly states have agreed to. Some states have done everything possible to shield themselves from any form of review mechanism. Some states, such as the United Kingdom, have submitted declarations accepting the compulsory jurisdiction of the International Court of Justice, for example, and other states, such as China and Russia, have not done so. Some states have added reservations to any treaty that has a dispute settlement clause, again, to avoid the possibility of going to a court or some kind of review mechanism. In some way, the question of enforcement is tied to what precisely has been agreed by a specific state.

That said, what we discussed prior to this, with decentralised enforcement, countermeasures and the different ways in which we can seek to induce compliance, is still relevant whether we are talking about China or any other state, because these are rights that we have as states under general international law.

The degree to which we will achieve meaningful compliance with the law will depend on the availability of procedural mechanisms. Why do we have a good sense as to what human rights law requires? It is because we have courts. We have tribunals. We have review mechanisms. There is a constant type of review of these treaty instruments, such that we have a good sense of what exactly they provide.

It is not the case with humanitarian law, but within frameworks such as the CCW we have a compliance mechanism whereby states have to submit information. There are many reporting obligations under different conventions, so to your question of enforcement it does not depend on having new rules. It is about trying to use every possible mechanism that we have, procedurally, to ensure that states know that they cannot get away with violations.

Q151 **Baroness Doocey:** What are your views on the creation of a dedicated domestic AI regulator with responsibility for AWS?

Tsvetelina Van Benthem: Without knowing more about the specific proposal and what this regulator would look like, I can say that in principle I see a benefit in having a regulator that looks only specifically at the challenges and risks that come from autonomous weapons. The questions I have are these. Where does it sit institutionally? What is its relationship with the Ministry of Defence and with the units that are tasked with legal reviews of new weapons? What is its relationship to this global regulatory mechanism that Prime Minister Sunak put on the table a few days ago? I presume that will look at broader questions that relate to artificial intelligence, but maybe some subsection of it would relate to the issues that we are discussing today.

Baroness Doocey: You would be in favour in principle, subject to working out how it would operate.

Tsvetelina Van Benthem: Yes, subject to many things.

General Sir Chris Deverell: A glib answer would be that I am not sure why we would need a regulator for something we do not intend to have. A more helpful answer would be that we clearly need controls in society to limit potential harms in AI generally, not just in AWS, while ensuring the delivery of many of its benefits. We clearly need that. The Prime Minister has committed to it. It is very high in the political agenda. It is a G7 summit topic, so I am sure we will have AI regulation. The question is what it would look like and, in particular, how it would balance the trade-offs between risks and opportunity.

Baroness Doocey: Part of this is because, as my colleague said earlier, we want to make sure that our report, when we do it, covers not just what is happening today but what might happen in the future.

General Sir Chris Deverell: You need to be careful that you do not set up a body for dealing with some very extreme use case and ignore the tsunami of really significant AI issues that need regulation. You could place the issue of regulating AWS somewhere underneath a broader AI regulator. That is an open question for me, but I certainly would not propose setting up a regulator specifically for AWS. That would be nonsensical.

Baroness Doocey: That is helpful.

Q152 **The Chair:** Sir Chris, you are on record as saying that 10% of defence budgets ought to be allocated to digital innovation. That was two or three years ago, but are you of the same view now? Do you think that shift in spending emphasis would cover some of the things we have been trying to tease out this morning?

General Sir Chris Deverell: Yes. The irony here is that we talk as if AI is everywhere in defence, when it is almost nowhere. The penetration of AI in the UK Ministry of Defence is almost zero. There are lots of reasons for that. We can talk about them, but the headline is that the gap between what the MoD says it is doing about AI and what it is actually doing about AI is vast, and by that I mean the delivery of systems into service. There is a lot of innovation theatre, experimentation and talking, but there are very few core programmes in defence that have AI in them.

That is partly for reasons specific to the Ministry of Defence and partly for wider reasons about the deployment of AI. It is quite easy to put AI into a point solution, an app on your phone. It does not disturb the rest of the ecosystem. It is an additional layer of something. AI often requires a system solution and a change of the whole organisational design and culture. Those are very difficult to do. That is a good reason for AI not penetrating. It is difficult to achieve those system solutions in any enterprise.

There are also bad reasons, to do with risk aversion, a reliance on a particular set of contractors that we have always relied on, and the general ills of procurement in the Ministry of Defence.

The Chair: Thank you. You have given us ammunition, so to speak, for when we see the responsible Minister later in our inquiry.

- Q153 **Lord Hamilton of Epsom:** Surely one of the other reasons for an independent voice on AI is to reassure public opinion? At the moment our beloved press loves emphasising the Armageddon aspects of AI—that everybody will lose their jobs and we will have machines liquidating vast numbers of people indiscriminately. We need to bring a bit of balance into this from somebody who knows how the whole thing works.

General Sir Chris Deverell: Yes, I totally agree. The hype is overtaking us at the moment. For me, the probability of AI representing an existential risk is very low—not zero, but very low. That is the climate into which we are currently operating. It is important to find ways to reassure the public that these probabilities can be limited or drastically reduced or controlled, and, at the same time, educate people on the massive potential benefits that exist, which are huge.

Lord Hamilton of Epsom: Would you agree, Tsvetelina?

Tsvetelina Van Benthem: Yes. In addition to the point about public trust, I was wondering about the specific modalities of what such a body may look like, because it would be good to have a body that is a little more inclusive and transparent about what exactly it is looking at. We have seen that different processes that are taking input from a variety of stakeholders are sending a good message, but of course this is subject to answering the key questions about what its relationship will be to other bodies that are already looking at the types of systems we have been discussing today. We do not want to have clashes and difficulties of co-ordination. As was already mentioned, there may be some mixed signalling already happening on questions related to autonomous weapons systems. We want to streamline the process rather than create even more places for obstacles along that way.

- Q154 **Lord Grocott:** I want to make a quick reference back to an earlier discussion, which is relevant to this question, about non-state actors. When you are talking about universal compliance, the possibility of getting international agreements and all the rest of it, as one or two of our witnesses have mentioned, it is not unhelpful to think for a moment about the non-proliferation treaty, which was never universally accepted. You could say that it was certainly accepted by all the key players. None the less, at least in my submission—you may disagree—it has been an extraordinarily effective international arms control agreement. Part of the reason for that relates to this specific question, in that there were advantages for non-nuclear states to observe the treaty. There were advantages of sharing knowledge about nuclear power and the rest of it.

This question is about the dreaded subject of non-state actors, which seems to get more and more significant. We have had a bit of evidence

on that so far. People have made the point, which you may agree with, that it is quite a serious issue, certainly in comparison with the use of nuclear weapons, because it is allegedly much easier for a non-state actor to develop an autonomous weapons system than it is to develop a nuclear weapons system.

You can argue the point on that, but in your judgment how easy is it for non-state actors to acquire this kind of defence instrument, given that there is no advantage to them in not doing so, as far as I can see, which is one of the mechanisms that made for the success of the NPT?

General Sir Chris Deverell: I think we could debate the effectiveness of the nuclear non-proliferation treaty which has not stopped certain states pursuing and developing nuclear weapons. But, on your broader point about non-state actors, this is an area that we need to think about and not just assume that the debate is always about state actors. As to whether non-state actors could acquire them, buy them, develop them or be given them, it depends a bit on what you mean by an autonomous weapons system.

Lord Grocott: Could they do so within your definition of it?

General Sir Chris Deverell: On that definition, they would be hard for non-state actors to develop. Whether any state would sponsor a non-state actor by giving it those kinds of weapons must be debatable, given that those weapons could then be used against that state. Some things are closer to the end of the spectrum. There are loitering munitions that could fall into the hands of non-state actors and could be very dangerous in those hands. They are hardly likely to represent the scale that would pose an enormous threat to the existence of humanity, but could they pose a threat to Downing Street? Yes. We definitely need to think about how we handle the issue of non-state actors acquiring these kinds of weapons systems and how we counter them.

Tsvetelina Van Benthem: I initially hesitated to answer, because empirically speaking I am not competent to say how easy it is for them to acquire these systems, but legally this ties to the discussion we had about other rules of international law that may have a bearing on the regulation of autonomous weapons systems. It may be worth thinking about different positive obligations with a due diligence standard, whereby states would have to take certain measures to prevent the acquisition of such weapons by non-state actors, whether this is under due diligence obligations in custom, meant to protect the interests of other states, or under human rights law, meant to protect the interests of individuals. There is an interesting dimension to this that is legal, which is to map out all the relevant obligations that would pertain to the potential acquisition of such weapons by non-state actors and the foreseeable risks that it poses to both states and individuals.

The Chair: The final question is especially designed to allow a one-sentence answer.

Q155 **Lord Fairfax of Cameron:** If each of you could make one

recommendation to the UK Government in this area, what would it be?

Tsvetelina Van Benthem: The cheeky way is to say that I will add two components to this. It is to lead and to learn. We have to lead on the law, as the UK has already put forward in the intergovernmental discussions, and go deeper than just saying that we have to clarify the law; otherwise this will become a tool for discursive procrastination. We then have to learn, because of the possibility of disappointment. We have to listen to others to see what their concerns are. The combination of these two will place the UK in a very good position in international discussions.

General Sir Chris Deverell: This is my single sentence, which I would like the right to expand upon slightly: "Don't throw the baby out with the bathwater". I am worried about the boundary between AI generally in its different applications and autonomous weapons systems. The fact that you are even having a conversation specifically about autonomous weapons systems is drawing attention to a case that is at the extreme end of the spectrum, when there is a lot of work to be done around AI generally. Most of it would be led in the civil sector, but it needs to then be carried into defence.

The risk is that one does things in the space of AWS that then make the employment of AI in defence exceedingly difficult. We absolutely have to avoid that. The Government have a duty to avoid putting the UK's armed forces at a major technical disadvantage to our likely adversaries, which would be the case if they have AI and we do not.

The Chair: Thank you very much. Tsvetelina Van Benthem and Sir Chris Deverell, it is to state the obvious to say that the subject of our inquiry is a complex one. This morning we have been dealing with one of the most complex areas, but we are all very, very grateful to both of you for helping to unravel some of the tangled skein of opinion, fact and conjecture that we have encountered so far. Thank you both very much indeed.