



OFFICIAL REPORT
AITHISG OIFIGEIL

Economy and Fair Work Committee

Wednesday 19 November 2025

Session 6



The Scottish Parliament
Pàrlamaid na h-Alba

© Parliamentary copyright. Scottish Parliamentary Corporate Body

Information on the Scottish Parliament's copyright policy can be found on the website - www.parliament.scot or by contacting Public Information on 0131 348 5000

Wednesday 19 November 2025

CONTENTS

	Col.
SUBORDINATE LEGISLATION	1
Public Procurement (Iraq Free Trade Agreement) (Miscellaneous Amendment) (Scotland) Regulations 2025 [Draft]	1
Cross-border Public Procurement (Miscellaneous Amendment) (Scotland) Regulations 2025 [Draft]	2
ARTIFICIAL INTELLIGENCE (ECONOMIC POTENTIAL)	6

ECONOMY AND FAIR WORK COMMITTEE

32nd Meeting 2025, Session 6

CONVENER

*Daniel Johnson (Edinburgh Southern) (Lab)

DEPUTY CONVENER

*Michelle Thomson (Falkirk East) (SNP)

COMMITTEE MEMBERS

Sarah Boyack (Lothian) (Lab)

*Willie Coffey (Kilmarnock and Irvine Valley) (SNP)

*Murdo Fraser (Mid Scotland and Fife) (Con)

Stephen Kerr (Central Scotland) (Con)

*Gordon MacDonald (Edinburgh Pentlands) (SNP)

Lorna Slater (Lothian) (Green)

*Kevin Stewart (Aberdeen Central) (SNP)

*attended

THE FOLLOWING ALSO PARTICIPATED:

Steve Aitken (Intelligent Plant)

Kayla-Megan Burns

Leo Fakhru (XYNQ)

Alasdair Hamilton (Scottish Government)

Dex Hunter-Torricke

Ivan McKee (Minister for Public Finance)

CLERK TO THE COMMITTEE

Anne Peat

LOCATION

The James Clerk Maxwell Room (CR4)

Scottish Parliament

Economy and Fair Work Committee

Wednesday 19 November 2025

[The Convener opened the meeting at 09:00]

Subordinate Legislation

Public Procurement (Iraq Free Trade Agreement) (Miscellaneous Amendment) (Scotland) Regulations 2025 [Draft]

The Convener (Daniel Johnson): Good morning, and welcome to the 32nd meeting in 2025 of the Economy and Fair Work Committee.

This morning, we will be considering two Scottish statutory instruments before taking further evidence as part of our work on artificial intelligence. To begin, I note apologies from Sarah Boyack, Lorna Slater and Stephen Kerr.

I welcome Ivan McKee, Minister for Public Finance, and—online—I welcome Alasdair Hamilton, procurement policy portfolio manager, and Jess Gray, lawyer, from the Scottish Government. They will present on the SSI, and I invite Ivan McKee to make a short opening statement.

The Minister for Public Finance (Ivan McKee): As committee members know, while the negotiation of international agreements is a reserved matter, implementation can occur in devolved areas. In particular, agreements often include provisions providing for reciprocal access to public procurement. Accordingly, Scottish procurement regulations set out that bidders from countries to which a relevant agreement applies are entitled to equal treatment when bidding for contracts in Scotland.

The instrument updates the list of relevant agreements by inserting a reference to a new agreement between the United Kingdom and Iraq. The effect will be to ensure that suppliers from Iraq will be entitled to the same treatment as Scottish or other UK suppliers when bidding for contracts that are covered by the agreement. Under the terms of the agreement, Scottish bidders will be entitled to equal treatment when bidding for contracts in Iraq.

It is worth noting that the agreement is based on the agreement between the European Union and Iraq, which ceased to apply to the UK after Brexit. In that sense, the instrument can be largely

regarded as a restoration of rights and obligations that existed until that point.

The Convener: I have a question for the minister. For clarity, does the instrument mean that Iraqi firms need to be given equal consideration when bidding for public work, and vice versa for Scottish firms in Iraq? In practical terms, looking at the balance of trade and the industrial and economic base of both countries, are there likely to be Iraqi firms that bid? Has there been any analysis of that?

Likewise, what might the opportunities be for Scottish firms to bid for public contracts in Iraq?

Ivan McKee: In both directions, the trade is currently pretty minimal. We would hope that that will grow over time, although Iraq is not one of the priority countries in our export plan.

Data that we have—this is from His Majesty's Revenue and Customs regional trade statistics—show that £45 million of goods were exported from Scotland to Iraq in 2024, but the data shows only £1 million of goods coming in the other way. Both of those are about 8 per cent of the UK total.

The Convener: I thank the minister for that answer about the economic value. Are there any practical or other considerations?

Ivan McKee: Officials can correct me if I am wrong, but I do not think that there are any instances of Iraqi firms bidding into Scottish public procurement as yet. This instrument is just to add them on to the list of countries that can do so.

The Convener: As members have no questions, we move to agenda item 2, which is the formal consideration of the motion. I remind everyone that only members and the minister may take part in this part of our proceedings. I invite the minister to speak to and move motion S6M-19302.

Motion moved,

That the Economy and Fair Work Committee recommends that the Public Procurement (Iraq Free Trade Agreement) (Miscellaneous Amendment) (Scotland) Regulations 2025 [draft] be approved.—[Ivan McKee]

Motion agreed to.

Cross-border Public Procurement (Miscellaneous Amendment) (Scotland) Regulations 2025 [Draft]

The Convener: We move now to agenda item 3, which is hearing evidence on cross-border public procurement regulations. I invite the minister to make a short opening statement.

Ivan McKee: It has become commonplace over recent years for contracting authorities, bound by public procurement legislation, to co-operate where it makes sense to do so. Until recently, it

was a fairly straightforward exercise as the procedural rules on awarding contracts—on such matters as minimum timescales and the type of notice that must be published—were substantially the same across the whole of the UK.

Earlier this year, however, the UK Government commenced the Procurement Act 2023, which had the effect of changing the procedural rules that are applicable elsewhere in the UK. That meant that new provisions were needed to cater for the scenario in which a devolved Scottish contracting authority is awarding a contract under an arrangement that is put in place elsewhere in the UK, or vice versa.

The Procurement Act 2023 gave both the UK and Scottish ministers powers to address the issue, and the Scottish Parliament agreed a legislative consent motion on the point. The UK Government exercised its power following engagement with us and, through public consultation, set out in legislation the elements of the 2023 act that will apply to devolved Scottish authorities when they are awarding a contract under UK arrangements.

This draft instrument seeks to mirror the provision made by the UK Government by setting out the provisions of the Scottish procurement legislation, which will apply to UK authorities when they are awarding a contract under devolved Scottish procurement arrangements.

Both the instrument and the UK equivalent legislation are drafted with the intention of applying only such of the provisions of the other regime as are necessary to allow the contract to be awarded. Those provisions can reasonably be described as procedural rules rather than policy-driven rules. The exception is where two or more authorities are conducting a joint procurement exercise and one is leading the exercise on behalf of the others. In that case, the applicable regime would be the home regime of the authority that is conducting the exercise.

This is a relatively technically complex instrument, but it is intended to allow co-operation on procurement to continue as appropriate.

The Convener: Thank you very much, minister.

I begin again by asking what the instrument will mean in practical terms. My understanding is that it will essentially enable devolved bodies to use UK-wide procurement frameworks. Are there practical examples in which either that has happened or the Scottish Government or its agencies intend to use the provisions? Can you outline some examples?

Ivan McKee: I will ask officials to come in on the specifics.

If we look across the public sector landscape, a huge amount of work has been done over recent years to encourage the use of framework contracts in procurement. We have saved hundreds of millions of pounds by doing that in the Scottish context, with different public bodies co-operating on those frameworks. The same can apply across the whole of the UK, where it makes sense to do so.

That has been the case up to now, and all that the instrument does, given the new UK procurement legislation, is enable the process to continue. It dovetails in the different procedural requirements so that co-operation continues to be technically possible.

I will ask officials to come in if they have any specific examples. If there is any more detail that you need, convener, I am happy to come back.

Alasdair Hamilton (Scottish Government): I can provide some specific examples of co-operation. In the sense of Scottish contracting authorities using UK-wide frameworks, Crown Commercial Services is the biggest provider of the big national frameworks on the likes of information technology. Advanced Procurement for Universities and Colleges in the higher and further education sectors in Scotland is a prime example of where co-operation flows the other way. Some frameworks put in place by Scottish HE and FE institutions are used by their counterparts across the rest of the UK.

The Convener: I have a subsequent question. Procurement is a topic that comes up regularly at this committee, and one of the regular points is about how open procurement processes are to the broadest range of firms, especially smaller firms. Will the instrument help, or is it largely to one side of that topic? I thought that it was important to raise that.

Ivan McKee: Absolutely—it is a very important topic.

Technically, the instrument is a continuation of what we have been doing. It puts in place the regulations that allow us to continue to do what we have done up to now, which opens up procurement more broadly. What does that mean in practical terms? Obviously, if we are putting out bigger contracts—which is what happens when they are amalgamated across a wider geography—we run the risk of there being bigger suppliers bidding for them.

There is a separate but very important focus of the Scottish procurement policy on supporting small and medium-sized enterprises to access public sector contracts. That is thoroughly embedded in the process—and I think with some success. We are now at the stage where 47 per cent of Scottish public sector procurement spend

went to SMEs in the last published data. That compares with a number of around 20 per cent across the whole of the UK.

We have made some significant progress, but there is always more work to do.

The Convener: As members have no questions, we will move on to agenda item 4, which is formal consideration of the SSI. I remind members that only members and the minister may take part in this section. I invite the minister to speak to and move motion S6M-19428.

Motion moved,

That the Economy and Fair Work Committee recommends that the Cross-Border Public Procurement (Miscellaneous Amendment) (Scotland) Regulations 2025 [draft] be approved.—[*Ivan McKee*]

Motion agreed to.

The Convener: A short report will be prepared and published, and I invite the committee to delegate report clearance to me as convener. Is that acceptable to the committee?

Members indicated agreement.

The Convener: I suspend the committee briefly in order to change witnesses.

09:12

Meeting suspended.

09:14

On resuming—

Artificial Intelligence (Economic Potential)

The Convener: Under agenda item 5 we will continue our evidence sessions on artificial intelligence. We are pleased to have two panels this morning, the first of which consists of Dex Hunter-Torricke, strategic communications adviser and former head of communications at SpaceX and current head of executive communications at Facebook, and Kayla-Megan Burns, tech founder and board member at the Royal Scottish National Orchestra, both of whom are attending online.

I would like to begin by asking you both whether you think we are getting it right with regard to how we understand artificial intelligence and the skills that we seek to instil in young people and the wider population. Much of the discussion is about losing jobs and workers being displaced, but I slightly shudder when my daughters come home from school telling me that they are being told that they must not use any AI whatsoever.

My sense is that we should be thinking about what we can use AI for. What are the right questions and the right ways to use it? How we can use AI to maximise our skills and knowledge and the expertise of the wider workforce? What should we be doing to give people the right skills to maximise the use of AI? Dex Hunter-Torricke, I noticed you nodding. Can I bring you in on that question?

Dex Hunter-Torricke: Yes. Thanks so much for inviting me to join. I should just clarify that I am no longer with Facebook; I am not currently working for it.

In terms of skills, you are absolutely right that we want to recognise that students and young people can be doing useful things with AI. We should be encouraging that and figuring out the right framework to manage as part of the education system.

Speaking candidly, there are probably very few students who are not using AI in all their homework assignments, whether or not that is something that schools are permitting or encouraging. It is important to rethink how we set homework and assess student performance in a world where AI tools are ubiquitous and young people are absolutely jumping with both feet into the AI future.

In terms of what those skills look like, part of it is thinking about how you integrate AI usage into all parts of the curriculum. There is a plethora of different tools that are useful across subjects, and

AI capabilities are rapidly growing more and more sophisticated. It is important that we rethink how we are establishing curricula, assessing student performance and getting folks to experiment with these things continuously.

It is also important to think about the broader societal context. How can we create the kind of environment in which young people will want to be able to adopt these skills and be encouraged to develop them? It is probably going to be by encouraging adoption of these skills and developing literacy in them at all age levels across society. That is going to be really important for a broader set of agenda items, which this committee is obviously focused on, such as how we can build a high-growth competitive economy for the future. That is going to require strong adoption and understanding of AI tools in all sectors of the economy, and that is something for which, to be honest, very few ecosystems around the world have figured out the right strategies. There needs to be a much larger conversation about how we drive that.

The Convener: Good. We will be interested in exploring the number of strands that you have laid out.

Kayla-Megan Burns, I am mindful that my deputy convener, Michelle Thomson, would like to talk more in depth about the arts, but can I ask you a similar question on skills? From an artistic point of view, what sorts of skills should we be thinking about? Are there as many possibilities as there are risks when you are considering the arts more generally?

Kayla-Megan Burns: With the arts, AI is quite a double-edged sword. For example, I know many people who had nothing to do with the arts prior to the availability of large-scale commercial AI tools—such as DALL-E and Midjourney—who have now jumped straight into the arts and are actively using AI to support their careers. They are making good livelihoods doing that, so there are definitely opportunities there. A really key skill set in that is experimentation and innovation—not being afraid to get in and about these things, get it wrong a few times and, by doing that, figure out how to get it right. That is definitely a huge part of it.

On the risk side, however, there are substantial setbacks. For example, AI art and digital design is a great area to look at. Just as I know many people who now make great livelihoods from making AI ads and things like that and are utilising AI tools to speed up their creative processes and bring things from conception straight through to reality in a very fast timescale, there is also a setback for real artists who have existed in the industry for a long time.

For example, previously, if you were going to set up a business, you might go to someone who does graphic design to sort out your logos and brandings and things like that. You would maybe spend some money on that—probably in the region of hundreds of pounds—and you would usually keep it local, and that would be good for the economy. Now, however, with access to tools such as Midjourney and DALL-E, you do not need to spend that money, and the money does not go to local artists with those skill sets. That is definitely a risk. Instead of going to local artists, that money is going into the pockets of OpenAI and tech giants like that. We know that that harms not only our local job markets, but our economy, because every pound that is spent locally can recirculate up to the value of £5, and that does not happen when that money goes to tech giants.

In the creative industries, we have seen similar risks and impacts happen before. Music was always the canary in the coal mine. For example, we saw the dawn of Napster and peer-to-peer file sharing, which resulted in music industry revenues taking a severe hit between 1999 and 2009, and the same thing happened when Spotify came along. Although Spotify was kind of the regulator for peer-to-peer file sharing, it still had a massive impact on the music industry.

We are now looking at that sort of thing with AI. Current projections estimate that 24 per cent of music creators' revenues are at risk by 2028 due to AI-generated content. That is the equivalent of about £8 billion of cumulative losses across five years, and that is escalating to more than £3 billion annually by 2028.

Currently, more than 33 per cent of songs that are uploaded daily to platforms such as Deezer are AI generated—that is a very recent statistic; I think that that report came out only this month—and that figure has tripled in the past 10 months, which shows rapid AI penetration of these markets. However, people do not like it, which is the really stark thing. For example, 55 per cent of UK adults express discomfort with the idea of accidentally consuming AI-generated music and 77 per cent believe that unattributed AI compositions amount to theft or unjust use.

At the same time, 97 per cent of listeners cannot reliably distinguish AI music from human-created music, and that disconnect amplifies the risk of cultural dilution and economic displacement. Our own RSNO has managed to bring Hollywood to Scotland by producing film scores, which is great, and we have had significant economic impact from that, which has been exciting. There has been £17.2 million in gross value added for Scotland this year, supporting more than 300 jobs and 500 freelancers, and I think that there have been wellbeing benefits worth £11.6 million.

While that is great, there is a big risk with AI with regard to film scores. If 97 per cent of people cannot tell the difference between AI and human-generated music, imagine how that number would change if you were to put music over a really impactful lightsabre scene. At that point, I think that we would hit nearly 100 per cent of people not being able to tell. That is a real risk for you.

The Convener: Great. I am tempted to ask a raft of follow-up questions, but I think I would probably be in danger of allegations of copyright theft from my fellow members, Michelle Thomson and Murdo Fraser, both of whom want to follow up on some of those points. I will hand over to Michelle and invite her to ask her questions.

Michelle Thomson (Falkirk East) (SNP): Good morning. It is an absolute privilege for us to get the benefit of some of your precious time this morning.

I want to come to Dex Hunter-Torricke first. Your hinterland is quite startling, and you have recently started working with the Treasury. Given the private sector career that you have had thus far, what is your perspective as someone who has come in and engaged with the public sector?

Our Scottish Government is working on an AI strategy and plan at the moment, and I suspect that the challenge that it faces is what to make a priority when everything feels as though it is a priority and when you yourself have said that AI integration is more than a technology. What advice would you give the Scottish Government?

Dex Hunter-Torricke: That is a great question—it is the trillion-dollar question. There is a lot in the strategy and the overall framework that is absolutely world class. It is a very sound overall framework. The challenge will be executing against that and doing it at the pace and scale that are necessary in order for you to be competitive.

The technology is maturing so fast now. In a way, many folks have still not fully adapted to what has happened with generative AI in the past three years. I have seen numbers from the Office for National Statistics that say that fewer than 20 per cent of Scottish businesses are actively using AI. The number is probably higher than that, because many employees are using AI in an unofficial capacity on their personal devices. However, in general, adoption rates are probably quite low.

The AI that is developed in the next two to three years will be vastly more sophisticated and transformative than what we have already seen. It is worth bearing in mind that it is not even three years since ChatGPT arrived—the three-year anniversary is next week. In that time, trillions of pounds of economic value have been added to the world economy. There is some pretty plausible analysis that shows that, if, for example, AI and its impact on US economic growth are taken out of

the picture, the US economy is growing by only about 0.1 per cent. In other words, pretty much all the economic growth in the US is being driven by AI. Of course, that is a very different ecosystem, which is highly dependent on silicon valley, but it illustrates the scale of value that is being created at the moment.

In the next two to three years, we will see the much quicker arrival of systems that will be supremely transformative for those who are able to jump in quickly, harness those capabilities and enable them to rethink their organisational strategies, their culture and their leadership models—in other words, the very human and analogue things that need to go with the technology, which I would say are more challenging to implement than the technology. We will have little time to adopt those systems before other ecosystems—competitive economies and other companies—will look to seize that advantage around the world. That very much intensified pace of competition and the entry of new challenges in established traditional industries should be of great concern to us and to policy makers in all ecosystems.

09:30

In the past year, Sam Altman, the chief executive officer of OpenAI, has made the famous prediction, which is the talk of the town in silicon valley, that, at that some point in the next year, we will start to see the first billion-dollar companies that have a single employee, as a result of their being backed up by large amounts of very sophisticated AI agents. We have not reached that point yet, but we are certainly not far off it. There are companies that have been around for only a few years, and have as few as a dozen employees, that have already hit the billion-dollar valuation mark. We will probably see more companies getting into that space.

The nature of AI is such that it offers extremely asymmetric advantages. What do I mean by that? I mean that a small company that is able to move quickly to embrace cutting-edge capabilities and to orient everything around taking advantage of those tools might be able to disrupt much larger, slower incumbents. That is a huge opportunity for small businesses. I tell folks, including Scottish small businesses—I spend a lot of time in Scotland—that there are no small businesses any more. In this era, the smallness is entirely in our minds. The same could be said for countries. In this environment, there are no small countries. A country that can gear the ecosystem towards being ultra-ambitious and that can figure out how to quickly integrate those capabilities into the public and private sectors will gain huge advantages in the competitive global landscape.

However, that also means that it will be highly vulnerable to being challenged in the same way, so fast moving is the ecosystem today.

Michelle Thomson: There was so much in that answer. Working on the basis that, almost regardless of what people do, it will already be too late, I get the sense from what you are saying that we should not get in the way of the disrupters who will manage to create sole-employee, billion-dollar companies. However, when it comes to the utilisation of AI in the public sector, trust is a much bigger consideration. In the context of some of the use cases that the public sector deals with, getting it wrong could have catastrophic consequences with regard not only to the data, but to society's trust in government and all that that entails.

I would appreciate your thoughts on that.

Dex Hunter-Torricke: Trust is absolutely critical. Some pretty good polling and analysis have shown that the vast majority of consumers—well over 50 per cent of the population—are pretty sceptical about or afraid of AI. The more people learn about AI, the more afraid they become of it. It is not an educational problem.

There is a fundamental problem of legitimacy and trust in technology, which is closely linked to the way that people perceive big tech and the many reputational and regulatory missteps that the industry has taken over the past decade. However, if we separate big tech from AI generally, which involves far more than big tech, we can see that lots of systems are being created by innovators from small companies, including pioneers based in Scotland. It is important to think about that.

We need to think about how we can be transparent about the ways in which AI is being used and how it is using people's data while also illuminating the really transformative applications of AI, such as how it can be used to deliver improved public services at a lower cost. It is important that we do that so that we can have a reasonable mainstream public conversation about what we want to use such technologies for.

There are all sorts of applications for different types of AI that very few people will think are extraordinarily controversial. There are things that add incremental but important value to the way in which we deliver services. They do not necessarily involve using a lot of very interesting data. A perfect example of that is the fact that, for years, cities all over Europe and worldwide have been using machine learning systems to optimise the flow of traffic through cities and to cut congestion by linking into the traffic light system and figuring out how to optimise the timing of the cameras so that traffic can be moved around. That is not a new capability; it is very established. When they

learn that there are perfectly well-tested, proven strategies for dramatically cutting congestion flows and decreasing the amount of time that car engines spend idling, which leads to improved air quality and helps with the meeting of environmental targets, most people will say, "Sign me up."

On the other hand, there are applications in the healthcare system, where, naturally, there is much greater sensitivity. There will not necessarily be a right approach and a wrong approach, but there could be lots of different approaches with a lot of grey. In those areas, we need to have a conversation about what people are comfortable with.

Michelle Thomson: There are a lot of follow-up questions that I could ask, but I want to bring in Kayla-Megan Burns.

Earlier, you mentioned some statistics. I know that some of them came from the report on the RSNO's economic impact, because we held an event on that last week in Parliament, but it would be useful to know, for the record, where the other statistics came from.

Kayla-Megan Burns: They came from a recent report by Deezer, which is a music streaming platform. That is where the stats on the number of AI-generated songs uploaded to such platforms came from. Previous reports brought together reports from different platforms about where those songs were shared. There are further reports specifically about consumer attitudes towards AI music. I am more than happy to share links to those reports after the meeting, if that would be useful.

Michelle Thomson: Thank you very much for that.

I think that the RSNO has been very leading in putting you on the board, given the kind of concerns that many creatives have about AI. It would be useful to flesh out which sector in the creative arts has the most concerns. The RSNO has done a tremendous amount. I have seen the uptake of its live performances by audiences. You correctly pointed out that it has done some marvellous stuff with recording, such as its recent recording of the music for "Nuremberg" at its film studio.

However, there is something about the authenticity of live music. How do you see AI being able to be integrated to enhance the service offering of a live orchestra such as the RSNO? In other words, what ideas have you brought to the board of the RSNO about how it might be able to get ahead of the game?

Kayla-Megan Burns: It is interesting that, whereas, in the past, it was possible to separate

the creative industries into the music industry, physical art and so on, in the AI era, that approach is becoming less and less useful. However, there are still differences in context. For example, with AI-generated music, you cannot really tell the difference, and that is having a massive impact. With AI-generated art, we can see the same kind of impact in a different field. Previously, music was the canary in the coal mine—as we saw with Napster and Spotify—but that is no longer the case. Now, the impacts are happening across the creative industries simultaneously.

As you mentioned, music is a little bit different when it comes to live performances. Nothing can replace humans in that respect. A live performance is a shared cultural experience, and I think that, in this era, we have an opportunity to generate a more diverse cultural experience. Right now, all the publicly available information has already been scraped by AI training models. All the high-quality data sets have been scraped. That material has been exhausted. There is no high-quality information that is publicly available online that has not yet been touched. That is all gone. Unfortunately, that includes all our artistic works. No creators have been compensated for their work being used in those models, which is quite significant across the board.

However, we have a real opportunity, which relates to the fact that what is most valuable for such systems is diverse training data: edge cases, extreme cases, unusual data and things that you would not normally come across. In Scotland, we are in a really fortunate position in that respect, because we have fantastic artistic institutions. On top of that, we have the Gaelic language, which cannot be overlooked, because it is a minority language on a global scale. That makes it incredibly useful from the point of view of training data, but it is also important culturally.

We are in a really interesting position right now, in which, by investing in our arts and culture, we can simultaneously invest in AI. Rather than saying, on the one hand, “We need to support local artists because AI can’t take over,” and, on the other, that we should be investing in AI, because it makes things more accessible and makes it possible to get things done more quickly, we can use those as complementary rather than opposing arguments. In Scotland, we are uniquely positioned to take that approach, which could have a really dramatic effect.

That approach also gives us other opportunities. For example, instead of competing head to head in areas such as infrastructure, which involves high capital expenditure, changing our energy systems and making multiyear commitments, investing in our arts now to get that training data for AI would diversify incomes for musicians and

artists, which would have a positive impact on wellbeing. The music industry, in particular, has been torn apart in recent years. In 2019, the median self-reported income for self-releasing artists was under £13,000 annually, with 47 per cent earning less than £10,000 per year. That has serious mental health impacts. Investing in our arts would not only enhance the arts and culture and our wellbeing but give us a unique relationship with AI training data, by producing the edge cases that I mentioned.

There are good examples of that being done elsewhere. For example, Ireland’s basic income for the arts scheme has shown fantastic economic returns. Although that has not been linked to AI, it provides a fantastic opportunity to do that. Over three years, that scheme supported more than 2,000 artists, and it has now been made permanent. That model is low risk but high reward.

Michelle Thomson: We could unpick so much in that; that is our challenge with such a short, sharp and focused series of sessions.

I would like to get reflections from both of you on another point. About 10 years ago, when I was in Westminster, we talked about AI in a session with a professor from the University of Cambridge. At that point, it seemed unbelievable how many base functions of lawyers and accountants were going to be taken over, although we know that to be true now.

When I asked that professor what skill set was going to inherit the earth, his answer was that the creatives will keep on creating no matter what, and they will harness the power of AI to endlessly create—and that will be the merging point. That has always stuck with me, and I would appreciate hearing your reflections on it. Do you believe that that is true? What does it mean for how we fundamentally shape the provisioning of everything, from a governance perspective?

I will first come back to Dex Hunter-Torricke and finish with Kayla-Megan Burns, and that will be me done, convener.

Dex Hunter-Torricke: I would agree with that philosophy. People assign intangible value to all sorts of things, which reflects a set of societal preferences and values. Every bit of research tells us that wine is stored better in bottles with screw caps yet, overwhelmingly, consumers prefer wine from bottles with a cork, because it feels like a more premium product, even though it generally tastes less good.

Right now, AI systems absolutely can generate compelling and high-quality creative work across any medium. In the next 12 to 18 months, somebody is likely to generate a full-length Hollywood-quality film on their personal computer using AI. Hollywood is worrying a lot about that,

but will people feel that such content is worthy of the same recognition as work that is created by human artists? I think not.

09:45

When I was travelling through Edinburgh airport a few weeks ago, I wandered into a gift shop, where I noticed a bunch of AI-generated art on sale as souvenirs on a stand. I felt really sad about that, because Scotland has infinite numbers of fantastic artists, but this garbage that was created in about five seconds using a piece of AI software did not at all reflect the creativity of the Scottish art scene. It was a bit of a wasted opportunity. What tourist would go in there and buy AI slop? The shop could have offered something valuable and given artists a platform to offer visitors something memorable to take home.

On the overall skills picture, I was based in silicon valley 10 years ago and I vividly remember a huge debate among my colleagues about what skills we should be encouraging young people to learn so that they were more resilient in the future economy. The consensus was to get young people to learn to code and push them towards science, technology, engineering and mathematics subjects, because a future that was increasingly shaped by technology would require a lot more people to have those skills. I pretty strongly disagreed, and I would like to think that my position has been strongly borne out with the passage of time.

Machines are very good at coding. AI systems that are coming from Google DeepMind, which is based in London, now outperform more than 93 per cent of programmers who are so good that they take part in competitive coding contests. We are on a trajectory towards the vast majority of code being generated by machines—probably autonomously, or almost entirely autonomously—in the next few years.

Given that, what should we be investing in? What is likely to make us more resilient for longer in a future with systems such as that? It is not just creativity; it is having talent that is genuinely interdisciplinary and multidomain in its expertise. The challenges of a future that is more interconnected and moving faster than ever will very much involve a bunch of different problems appearing all at once. Companies that are at the forefront of adapting their culture, organisations and leadership to AI are starting to recognise that the future will probably belong to a new set of business leaders and managers who do not belong within existing narrow bands of roles.

Now we have marketers, public policy experts, project managers, legal experts and so on but, 10 or 15 years from now, there will probably be a

whole bunch of folk who are operating across all those disciplines at the same time—people whose skill sets are multifaceted and who can solve problems across all those areas very quickly, because they are backed up by world-class AI systems, which will be the best experts in all those categories.

We therefore want people who are genuinely creative—I mean that not in a narrowly artistic way but in a way that means that they can figure out how to solve problems while working across all areas very quickly, with a lot of context that probably does not confine itself to a single traditional academic subject area or professional domain.

Michelle Thomson: Thank you. I put the same question to Kayla-Megan Burns.

Kayla-Megan Burns: I could not agree more with what Dex Hunter-Torricke said. To be honest, I find the situation quite comical. Over the past decade, there has been a real undermining of the value of the arts and creativity, in favour of what are seen as “real” studies and “real” jobs, such as coding and hard STEM subjects. Ironically, because of AI, coding jobs are being cut, while demand is increasing for philosophers, because we are in an era when those kinds of skills are really needed.

Ethics will be a huge area, because what organisations such as OpenAI and Anthropic are doing means that we need to look at creating guardrails for how we use AI, what it is allowed to do and how we interact with it. What is considered “safe” and “right” is being decided by just a few people in these tech giants. That is happening on a global scale, because the tools are being used globally. Decisions about “right” and “wrong”, about what is safe or unsafe, and about what is allowed or not allowed are being made completely independent of the cultural contexts in which the tools are being used. I am very conscious that there is real space for addressing those issues now.

As for skills, I am so conscious—Dex Hunter-Torricke mentioned this—of young people being funnelled down a narrow track, specialising too early and being popped into narrow fields with people saying, for example, “You’re going to do computer programming—that’s going to be your thing.” Concern and disillusionment exist among young people because they have been told that the track to take is to do well in school, get into university, find a good job and climb their way up the ladder, and that they would do best by doing STEM subjects and working in such areas. If anything has been proven by this new AI era, it is that that paradigm is totally and utterly false. We see so many young people who are being offered inappropriate jobs or who are severely

underemployed or unemployed. Jobs are being cut severely—especially entry-level jobs, because they are the easiest to get AI to do.

That will leave us with a massive skills gap of people who are experienced and can, for example, quality check the work that AI systems do. If an AI system is autonomously coding, a qualified human can check the code, ensure that it runs correctly and fix any issues. That role will tend to be for someone who is senior in an organisation, rather than for graduates.

What is happening with grad jobs? We need much closer links with industry, rather than the tight and narrow track from education into a grad job. That model is clearly not fit for purpose now, never mind for what we will go into in the next few years.

Michelle Thomson: Thank you. I will hand back to the convener.

The Convener: First, as somebody who graduated with a degree in philosophy 26 years ago, I say thank you very much to Kayla-Megan Burns for validating my educational choices.

I will ask Dex Hunter-Torricke a brief supplementary question. I am interested in the notion that AI tips economies of scale on their head. How far do you take the points that you set out? In 20 years' time, to what extent will organisations be just one person configuring AI tools around them? How far will that go? I absolutely accept that you will see businesses like that, but will all businesses be like that? What will a sensible organisation look like in size and—more critically—in configuration? Will that be about how well you specify things? If we take the coding example, to get good code out of the AI, you still need to give it the right specification. Is that what the core function of an organisation will be? How far will this go? What will the functions be at the heart of organisations that seek to use AI?

Dex Hunter-Torricke: It is unsatisfying to say that we do not know how far this will go. In general, we can expect that, across the board, organisations will become much leaner in terms of human employees. Companies that require the largest workforces now tend to be doing things with a lot of physical infrastructure. Logistics, the supply chain and manufacturing are all areas where we are seeing huge advancements in robotics, so in warehouses where hundreds of people were employed, numbers are dropping quickly. In some cases, facilities are almost fully automated, and that is certainly where the future is going.

We are seeing movement in a large number of organisations—including big employers in the UK—to begin reducing the size of new talent pipelines. We have seen double-digit falls in

certain roles being recruited to, such as new graduate roles in consultancy, accountancy and other professional avenues that were previously quite stable.

We do not fully know how far that will go. There will still be value in having scale in particular areas, but for knowledge work, where we are not necessarily relying on having a lot of infrastructure, we might end up with really condensed organisations that have a dramatically outsized impact. A lot of companies have yet to fully fathom how transformative that might be.

The examples that were given of the impact on artistic and creative jobs were spot on. A lot of small agencies are shedding a lot of their workers because, to be honest, one or two people backed by a lot of AI can probably do what 10 or 20 people would have previously been required to do. As the transition unfolds, it will be extraordinarily challenging.

The Convener: I will hand over to Murdo Fraser.

Murdo Fraser (Mid Scotland and Fife) (Con): Thank you, convener, and good morning to the witnesses.

Dex, the point that you just made is what I was going to ask about. What does AI mean for the workforce? Last week, we were looking at a report from Microsoft about the sorts of jobs that might suffer from development of AI, and in the top five were writers and authors. What does that mean for human creativity? What will the role be in future for original, human-created output? Is AI effectively just derivative on the work humans have done? If we are squeezing humans out of the picture, what does that mean? Does it mean that we will not have innovation in the future?

Dex Hunter-Torricke: First, AI is not just derivative. That might once have been true, based on the sort of crude AI systems that existed maybe five or 10 years ago—which sounds quite recent but is a lifetime in the industry. Now, we are in a moment where AI systems are generating things that, combined with human expertise, are transformative.

A perfect example is that we are experiencing the largest boom in scientific discovery in history. There is compelling research that shows exactly how the impact of generative AI systems is turbocharging the ability of scientists to do their work. I have looked at research that has been peer reviewed that shows that scientists who use fully generative AI-optimised workflows are publishing about 60 per cent more papers every year than their colleagues, they are getting promoted faster and they are getting three times more citations. We are seeing a huge explosion of

different discoveries across any number of domains.

Where does new technology come from? It starts in the lab and comes from science, and then it gets commercialised and becomes something that impacts our day-to-day lives. What is happening would not be possible without AI. People on their own do not have the cognitive power to deliver in the same way at the same scale.

However, AI will be extraordinarily challenging for societies, and you hit the nail on the head with the question. What exactly will be the impact on people's creativity, what will be the impact on people generally, and what are people for? There is a bunch of unanswered questions for a future that is not that far off—I do not think that it is 20 years away, unfortunately.

All the world's most valuable technology companies today—the trillion-plus juggernauts—are officially committed 100 per cent towards achieving AGI, or artificial general intelligence, which is essentially human-level intelligence machines. We are not anywhere close to that now: all the systems we have right now are quite crude in comparison. However, almost all the major AI lab leaders believe that we will reach AGI in the next decade, and some believe that we will achieve it much sooner.

There is therefore a world coming where the systems might end up excelling not just in the arts and creativity but across the expanse of the economy, and at least matching or potentially vastly superseding the performance of people. That is something that no society is prepared for at all.

We need to begin taking the issues seriously. If we want growth 10 or 20 years from now, which is no time at all, it requires us to be doing a whole bunch of things right now to prepare for that future. This change will be dramatically greater and more challenging than any previous technological transition that our societies have faced.

10:00

Murdo Fraser: You are probably not helping through giving us answers, but you are maybe helping us to ask the right questions. That is progress, so thank you for that.

I turn to Kayla-Megan with a similar question, but perhaps put it more in the context of music. Do you have similar concerns about how we create original music in the future if AI will just do it better?

Kayla-Megan Burns: I am going to come in with a very controversial statement.

Murdo Fraser: Good.

Kayla-Megan Burns: I do not believe that we have been creating a significant amount of original music for, at this point, decades, because of commercial trends. The question has been studied in multiple papers and, again, I am more than happy to provide them after the meeting.

Multiple studies have shown that our music has become simpler, more repetitive and, in some cases, more violent. It has become much simpler and more repetitive because of commercial pressures. To make the number 1 hit, or something that will get on the radio—and perhaps make some money in an industry in which people struggle to survive—the commercial pressures have limited creativity over the past few decades. That has been documented thoroughly.

I believe that we are on the edge of a creative renaissance and that AI will help us to achieve that. I think that it will change our entire perception of the arts and of value within the arts. As I have mentioned before, the market is shifting from mass scraping to licensed and curated data sets. Major deals have happened between AI firms and music majors—such as between Suno and Sony and between Udio and Universal—when previously there were litigation cases for unfair use from scraping of the music data sets. When I say data sets, I mean art, music and creative achievements.

There was an interesting ruling in the German courts recently—I do not know whether you will have seen it. It found that OpenAI's ChatGPT had illegally harvested copyrighted song lyrics, affirming that online creative works are protected by copyright law and cannot be freely used for AI training without permission. What is really interesting is that that has already happened: as I mentioned previously, all publicly available works have already been scraped. The ruling is a really interesting precedent to set because it means that, in effect, companies have violated copyright by scraping those works. It also means that it will apply across practically all AI companies, which have used the scraping strategy rather than ethically curated data sets.

The ruling also means that diverse edge cases and culturally rich data sets do not just add value for AI models to improve output quality and inclusivity; they are also really valuable for our economy, our people and our arts. We have seen some interesting cases. I forget the name of the firm—it might be Bronze, but I will get it to you after the meeting—but creatives in London used AI in combination with some artists to create music that changes every time you listen to it. It is a really experiential piece, and it is completely different from what we typically experience.

We are going to start seeing a lot more of those kinds of projects. Anyone can produce anything of relatively good quality very easily with AI, in music, in art and, to be honest, in most things—even apps and so on. What is going to stick out and be attention grabbing will be unusual and unique things that AI would not achieve alone and which involve elements of human creativity. That is where our strengths are going to be.

We have seen something similar before, albeit on a smaller scale, with peer-to-peer file sharing. All sorts of things were said about file sharing at the time, including that it was going to kill the music industry because no one would get paid for creating art. That did not happen. We now have artists who are entirely broke, living on very little money, but they are still actively creating. I do not think that anything will stop humans creating—it is part of our human nature. The question is how we protect that creativity and do that in an ethical and sustainable way. That will be a really challenging process, but I hope that a creative renaissance will be part of it.

Murdo Fraser: That is quite an optimistic outlook, which is good to hear.

I will follow up on one point you made about protecting intellectual property, which I think is an interesting one for us to look at. I can ask AI to produce me a piece of music in the style of, say, Beethoven, and it will do that. Beethoven is long dead and long out of copyright, so there are no IP issues.

If I ask AI to produce me a piece of music in the style of, say, Lewis Capaldi, it will do that too. However, Lewis Capaldi is still with us, he is still producing music and his music is protected. How does Lewis Capaldi protect his brand when anybody can produce a song that sounds just like him?

Kayla-Megan Burns: Multiple steps are needed, and I do not think that anything that has been proposed so far takes enough steps.

Lewis Capaldi is a fantastic example. Obviously, he is wonderful, he is Scottish, and it is great that we have him. However, there are many artists in Scotland who are not Lewis Capaldi, and they are being sorely missed in all of the arguments, even though they are a vital part of our arts and our creative industry.

Protections need to be in place across the board, regardless of the arts: protections for our likeness, our face and our voice. All those things need to be protected. That is not just from an artistic perspective, but also from a security perspective, so that no one can just create an AI avatar of me and come along to this committee and say things in my voice that I would not say, or represent my views in a way that I would not

represent them. That is important, not just in the arts but across the level.

In art specifically, we need various things. One of the first things we need is transparency that enables identification whenever works are used. Right now, the European Union requires AI companies to summarise training data by category. The UK is considering a similar approach, but neither of those approaches actually helps independent creators to identify whether their works were used in the data sets. The requirement is too vague; it is too broad. We cannot look at the information and say, “That’s me in there—that is my art. I was used”.

We need binding requirements for AI developers to maintain searchable work-by-work registries or, at minimum, provide audit access to creators who suspect that their work was included. Without that, opt-outs are meaningless because creators cannot identify what they are opting out of, and licensing therefore becomes impossible.

Secondly, we need a centralised licensing infrastructure designed for individual creators, not just intermediaries. The music sector has seen licensing deals between major labels and AI companies—companies such as Sony and Universal Music Group. They have started happening, and they protect the likes of Lewis Capaldi, Taylor Swift and Ed Sheeran. Those deals are being made on a large scale.

Unfortunately, however, our independent creators are being left out in the cold to fend for themselves, which is not remotely feasible. We need a statutory requirement that any licensing revenue negotiated through industry bodies includes a distribution mechanism for the independent creators—one that does not require formal collective membership, such as through a label. That could be a statutory licensing pool where independent creators opt in and receive allocations proportional to their content’s use in training.

The third thing that we need is an enforcement capacity that does not require individual litigation, because that is what we are stuck with right now. If an independent creator discovers that their work was scraped without authorisation, they currently have to fund their own lawsuit against well-resourced AI companies, which is an aggressive uphill battle. People are just not capable of doing that. Therefore, we should establish a statutory right to small claims and copyright adjudication for infringement claims under, for example, £50,000, with UK courts empowered to award attorney fees to prevailing creators. That would shift the enforcement burden from individual creators to the legal system itself.

It is important to recognise that the frameworks—the three things that I have mentioned, as well as the protection for voice, image and so on—would not apply retrospectively. We will still have the massive issue of the artists whose portfolios were used to train ChatGPT and other AI models in the past three, five and 10 years, because there is currently no recourse. Therefore, something like a one-time remediation mechanism—perhaps a statutory fund financed by AI companies for retroactive use to establish historical use matters and not just future protections—would also be important, because there is a massive gap right now.

Murdo Fraser: You have given us a lot to think about, and some helpful ideas about what changes need to be made in a policy perspective to protect original content. That was very useful—thank you.

The Convener: Thank you. The deputy convener would like to ask a brief supplementary.

Michelle Thomson: Yes. It is just a tiny point, which I do not want to take too much time on, but is the issue not even more complex than that with music? As you have explained, everything has been scraped, but you can create entirely new pieces made up of the best of the rest, if you like. I could sit and listen to Mahler 5, for example, and I could tell you which player it is in the trumpet solo in the opening; I could listen to “Nessun Dorma” and tell you whether the tenor singing the top C is Pavarotti, Domingo or Kaufmann. You could basically splice the best of the rest. It is not as simple, surely, as just taking an artist in a song or whatever; you could create something note by note with key thematics.

Kayla-Megan Burns: Yes, 100 per cent, and that is why it is important for AI companies to have databases and track what is in there and how it is being used. For example, if 0.8 per cent of a guitar piece or a rhythmic thing used in a Taylor Swift song is being used in AI-generated output, we need systems that can track the use of that—not only what is going into it, but how it is being used in AI systems. We are seeing interesting developments in those areas. For example, we have the C2PA—the Coalition for Content Provenance and Authenticity—which is a collaborative effort led by Adobe, but also endorsed by the likes of Google and Apple, to track who owns what. Right now, we have YouTube IDs, which do what it says on the tin. Essentially, this would be assigning something like a YouTube ID, in an invisible way, to every single piece of art that exists so that you can track how it is being used, where it appears in these systems, and proportionally how much of it is being used. It is important that we are able to do that, so that

appropriate compensation and attribution can be assigned.

That is a tricky technical issue, but we are seeing a lot of progress on that front. Using the technical effort that is required as an excuse not to do those things is a relatively poor effort, considering the era that we are in. We are seeing a lot of those developments. The fact that it is challenging does not mean that it should not be done, especially for the state of our arts, for appropriate attribution and for the wellbeing of our culture.

Gordon MacDonald (Edinburgh Pentlands) (SNP): Good morning. We have talked a lot this morning about copyright and AI, but there are also concerns from the public about misinformation, built-in bias, privacy violation and data harvesting. Furthermore, given the economic situation that Dex Hunter-Torricke outlined earlier and the swathe of jobs that could be undermined by AI, there will be a huge economic impact in every country if this comes along. That is not to mention the fact that the International Energy Agency has identified that energy consumption for data centres could be 1.5 per cent of global electricity consumption. Given that Governments are not very fleet of foot—we have had AI for 30 years, but ChatGPT has been around only since 30 November 2022—what regulations need to be brought in, and what should be the focus of any Government regulation?

Dex Hunter-Torricke: You are probably going to need vast amounts of distinct types of regulation, including national and global architectures for managing a bunch of these problems. All these technologies and the nature of most of the challenges are international. They are not things that most countries would be able to manage through any combination of purely domestic levers.

10:15

Among the broad buckets in which you need to think about having a public framework for responding, absolutely there is the economic disruption that will unfold, given the nature of how these systems will impact existing jobs and industries. There is also the geopolitical environment and the fact that we probably do not have an international framework anywhere near as sufficient as is required to figure out some of the common solutions that we need for managing these challenges.

You talked about disinformation. Many of the disinformation challenges that we are facing are taking place on platforms that are controlled by American companies operating in very different regulatory frameworks and with very different

political impulses. There is no solution here unless multiple countries and Governments can work together to tackle these problems in a much larger way. Simply put, most countries are not going to have leverage against the platforms to lead to meaningful change in the way in which they approach content decisions.

You have a bucketful of challenges with resource consumption and energy use, and those things are also part of the global fight against climate change, which is going very poorly indeed. A lot of attention is paid to the resource consumption of AI itself; the IEA stats are spot on, and those things are very alarming. A bigger issue that not enough attention has been paid to is simply that every conventional thesis about the value of AI is that it might allow our economies to grow. That is the promise of it and that is why you are seeing the hundreds of billions of dollars being invested in training these systems and in the vast infrastructure. If that happens, resource consumption by our societies will continue to grow. At the pace at which we are consuming resources with the kinds of environmental impacts involved, that is unlikely to lead to a situation where we can meet our climate change goals. In fact, we will get further away from that and could end up in a situation where we destabilise parts of the earth's climate systems that are already very frail.

Then you have the bucket of societal challenges. Everyone in the room is probably thinking through the nature of the migration crisis. We are in a world where, potentially, in the coming decades, we will see a lot of different countries and economies becoming increasingly uncompetitive in the face of the ecosystems that are able to optimise for this future with these kinds of technological tools. The ecosystems that are no longer competitive will not just see plateauing economic growth; they are likely to see systems that are in full retreat, with a corresponding rise in political populism and an overall model for societies that may not be workable.

What does the regulation look like? There are a bunch of individual strands that we can look at in any of those buckets, but, broadly, we need to think about those categories and recognise what it is realistic for us to drive ourselves within our jurisdiction, as well as where we should simply put our hands up and say that it is something that we need to have a real conversation about, involving leaders from multiple Governments—and quickly, because right now there is a large missing conversation on a number of those pieces.

Gordon MacDonald: On the point about multiple Governments having that conversation, the UK is no longer part of the EU. Are there any jurisdictions or organisations such as the EU that

are looking at this issue and have legislation in play that would be helpful for us to learn from?

Dex Hunter-Torricke: Multiple Governments in the EU have been looking at these issues and doing a lot of things that are very impressive. Again, we are probably missing a much larger holistic set of conversations that knit together action across the big domains that I talked about. Take countries such as Estonia, with its leadership in e-government and investments in technological literacy. That is very impressive and it has allowed Estonia to build up a position of global leadership in the way in which it is harnessing cutting-edge technologies to deliver citizen services. When it comes to AI, there are very hungry, innovative ecosystems around the world. South Korea has been deploying technologies for many years now that we do not have even now in Scotland or the rest of the UK. It is absolutely worth looking at how a number of other ecosystems are taking elements of these tools and transforming parts of their systems to take AI into account.

In the past few months, South Korea has passed probably the most ambitious set of education transformation proposals in the world for the AI era. It is a major investment—I believe that it is over \$700 million—into the education system to equip every student with AI-powered textbooks to provide true personalised learning at scale for the AI era. I believe that that will be simply par for the course for any world-class education system in the next few years. Those are things that we should be looking at.

Gordon MacDonald: Thanks very much.

I have a question for Kayla-Megan Burns, who said that there is a requirement for guardrails and that we need protections in place. Could you elaborate on what you see as the role of Government in regulation, given the information that we have received from Dex Hunter-Torricke?

Kayla-Megan Burns: I was speaking to people from Anthropic about this last night. Currently, we do not understand how AI models work. They are black boxes and we do not know how they make decisions. Therefore, how do we know that they are safe? How do we know what guardrails are needed, and who gets to decide what is safe, what is right and what is wrong? Currently those decisions primarily just sit with the tech companies, and that is an unreasonable concentration of power, considering how widely these devices are used globally.

We have spoken a lot about consumer-facing AI, but I am conscious that we are talking only about what is already with us. We are missing the horizon of what it will look like in the coming years. Right now, consumer-facing AI is confined to chatbots, but leading AI companies have indicated

that they will be moving away from that and bringing AI not just into smart devices but directly off screens and on to our bodies, for example, and into our surrounding environments. We will be entering an era of ambient computing, but that is not yet in any of the global conversations in the way it should be.

When you consider the impact that things such as our smartphones have had, particularly on young people and more vulnerable populations, that should be particularly alarming, because we do not know how these systems work. There are dark patterns in the ways in which we can be manipulated without knowing it, so we need to think about the guardrails and the safety of how we operate around that. We are not at risk of replicating the screen addiction that we have with internet scrolling, social media and smartphones as they currently exist. We are at risk of inventing something much worse, so I am very conscious that we need to get ahead of that, and that discussion is not currently on the table.

These devices could adapt not just to our schedule, our environment and our routines but to our neurotypes and how we process information. There could be a fantastic opportunity here if it is used right—for example, to democratise access to information. Whether you are a dyslexic banker or a bibliophile who is petrified by numbers, you should be able to access the exact same information to the exact same standard, because it will be able to adapt to how you process information. It will be able to present in formats that are the most appropriate for you to consume. As Dex Hunter-Torricke mentioned, there is a fantastic opportunity for learning through personalised textbooks and things like that—that is amazing—as well as increasing accessibility.

There is fantastic potential there, but there is also a serious risk of exploitation of, for example, biometric data, particularly things such as unconscious eye movements. I am sure that you are fully aware of things such as smart glasses, which have recently come into play. They have cameras on the inside that track your eye movements. Eye movements are a subconscious thing. We might not necessarily be aware of exactly what our eyes are looking at, but the cameras on the inside are able to track our eye movements, correlate to what we are looking at on the outside and potentially use that against us—for example by advertising against our subconscious, which is an invasion of privacy at an entirely undiscovered level.

Currently, that is not allowed. However, I have been speaking recently to people at Meta who are working in these areas, and the current processes for these types of innovation involve program typing, building use cases for it and then going to

market. There are no pauses to consider what the potential ethical implications are, what this looks like and how it will impact in day-to-day use. That is not there, and I think that it needs to have a much more significant role.

That said, I am very conscious that, particularly in the Scottish ecosystem, we tend to be very risk averse and are fond of setting up barriers to give a false sense of security. We need to be careful of walking that line. We do not want to inhibit innovation. We want to make sure that we can make the best use of the new innovations that are coming along, because there are fantastic applications, but at the same time we need to be conscious of the potential implications.

I would have looked ridiculous if I had come to you five or 10 years ago and said, “We need to look at copyright for the arts and we need to make sure that, if anyone scrapes information online, that can be tracked and people can be compensated for it.” If I had said that to the committee even five or seven years ago, people generally would have been looking at me and saying, “That is very unrealistic. We don’t need to do that.” Now I think that we are sitting in a similar situation with biometric data and AI as these devices move from screens on to skin, and I am aware of that. We are talking about personal implications at personal and societal levels on a scale that could put the Cambridge Analytica scandal, for example, to shame.

This is not scaremongering. I am an AI optimistic, I am a big believer in AI and I think that it can do fantastic things. I do not think that we should be shying away from it. We are not there with those issues yet, but they are 100 per cent on the horizon and they are not being brought into the conversation as much as they should be. It is important that we tackle the problems before they become widespread and prolific.

Gordon MacDonald: Thanks.

Willie Coffey (Kilmarnock and Irvine Valley) (SNP): Good morning. I invite you to say a few words each on the ethical side of all of this. Kayla-Megan Burns, you have mentioned ethics a few times, and colleagues have raised a number of issues that take us in that direction. Will you give us your thoughts on how we protect ourselves and society and also instil within the AI revolution a sense of responsibility, ethical behaviour and so on, or do you think that it is destined to just run its own course, in its own direction and at its own pace?

Kayla-Megan Burns: I stress that I do not want to see hindered innovation. We should not be setting up guardrails or extreme processes that will delay or hinder innovation. The AI era is here, and the next era of ambient computing is

thoroughly on its way—I expect that we will start to see aspects of that within the next year or so—but it is really important that it is done for a common good.

We will need to change a lot within our systems to achieve that—for example, by looking at what we commercially reward. Will we just let those who scrape the most data be the winners because they did that first, before we knew that it was an ethical problem, and therefore let them proceed unhindered, or will we reward ethical behaviours? There need to be commercial mechanisms to do that because, right now, industry's approach is, "Let's do this as fast as we can. We need to be the first and we need to be the most innovative. It doesn't really matter how we get there. We are not going to do the ethical testing first, because we need to see whether people will use it. We need to get into the markets before that happens." That means that we are waiting for the next scandal or for the development of issues such as screen addictions, so it has an important role in our societies.

10:30

The issue of neurodiversity comes up a lot, having come into the public zeitgeist within the past few years. We are much more aware of neurodiversity now than we were even five years ago. We can hypothesise from our understanding of neurodiverse populations and how they interact with technology that they tend to exhibit extreme behaviours faster. They could be our lead users for a lot of this. Rather than treating neurodiverse populations as vulnerable populations that need to be protected, we can look at them as having a lot of experience and knowledge due to sensory sensitivities or differences in information processing. For example, although screen addiction and the mental health implications of technology exhibit differently in neurodiverse populations, we see those impacts spread out across the broad population as well, but just on a different scale. Therefore, should we be looking at using neurodiverse populations as lead users and empowering those populations rather than exploiting them, which is currently a substantial issue in our society?

We will need to consider a lot of those big questions and to substantially change how we operate commercial models and what we value in order to make this ethical. Right now, it is a case of how we can make the most money. Recently, we have seen changes in how we engage with an AI model. Rather than asking it a question and it giving an answer back, it will often ask you five questions back. It might give you some information, but then it will ask you five questions back. That changed recently because it gets you

to spend more tokens, which is literally about putting more money into a machine. Is that in our best interests? No, because that gets us to spend more time in those models.

We have seen unhealthy use of those models and more time being spent on them, which is leading to issues, including, at their most extreme, AI-induced psychosis. However, it is more valuable for those AI models to get users to spend more time on them even though we know that that is unhealthy and not beneficial in its current form. We need to change a lot of those commercial paradigms and drivers to make sure that people make ethical decisions.

Willie Coffey: Thanks for that. Dex Hunter-Torricke, how can we throw an ethical blanket around this whole thing? Is it impossible or is it yet still possible? If so, who should do it?

Dex Hunter-Torricke: I think that it is still possible. It will require a whole bunch of leaders from the public and private sectors and from civil society to be having a mainstream conversation about the hard ethical choices. Right now, a lot of the conversation is in rooms like this one, with a small number of experts and leaders who are examining the issues in a detailed way. A lot of the ethical choices will be on profound matters, and there will not necessarily be a clear right or wrong answer; however, they will be choices that societies need to make.

Some of those questions are thorny. Generally, it will be really important not to frame those debates as being mostly about technical problems with technical solutions. One of the great missteps in a lot of the conversations about AI ethics over the past 10 or 20 years has been to talk about it literally as AI ethics. It is not; it is just ethics. It is a conversation about values, morality and what responsible leadership looks like for organisations. There is no clever algorithmic fix for a bunch of those things. The choices are down to business leaders and public sector leaders.

A perfect example is that, in the next decade, we will absolutely end up in a world in which we begin to see huge breakthroughs in medicine, using AI that allow us to tackle major diseases. You will see new cures for cancers that have been shaped by AI and there will be wonder drugs and therapies that potentially are life changing. Will the national health service be able to afford those? Will we be able to ensure that access to those drugs is available to our entire population, or will only a very tiny sliver of the population globally be able to afford drugs that give people years of life, with everyone else having to make do? That is a profound ethical problem, and it then begs a whole bunch of other questions about how we resource the healthcare system to make those advances available to everyone.

Another perfect example is the resource consumption that we have talked about a lot. In a world in which commoditised AI systems will allow any company to do amazing things very quickly, there are real resource usage concerns. Thirty seconds of generative AI video consumes as much power as running a microwave for an hour. Should we be creating apps and offering services that are just generating garbage when that will cost our society in some other way? Companies will have to think about those choices and then create a value system for and be transparent about them.

Willie Coffey: How do we persuade the single-billionaire company that you mentioned earlier to embrace this and to observe the ethical standards that we might want to deploy across the AI sector? How do we persuade that single-person company to do that?

Dex Hunter-Torricke: You have to start by talking about it much more loudly. Right now, the conversation with those companies and a bunch of the leaders of those companies is very narrowly scoped around the question, “How can I just acquire as much of your product as fast as possible to drive economic growth?” The nature of that growth and whether it is based on a real sustainable foundation—I mean not just environmentally, but whether it is built on a strong set of pillars of societal support—will be absolutely critical so that it delivers what we want for our societies over the coming years. Currently, that conversation is not mainstream; we are not hearing that from leaders.

Willie Coffey: Thank you, both.

Kevin Stewart (Aberdeen Central) (SNP): Good morning. This is the third week of our AI investigation and, I must be honest, there have been ups and downs in the evidence that we have heard. There may be huge positives and benefits from the AI revolution but, at the same time, we have heard that there are a lot of worries. I am sitting here thinking about what the masses of people at home who are watching this committee will be thinking. I am being quite sarcastic in saying “masses”, but these things create worries. We have heard about fully automated industries, and billion-dollar companies run by one person. We have heard about all the changes that could take place because of AI and that may make people redundant—some would say in more ways than one. What are the positives for those folk who may be sitting at home thinking, “Where do I fit into all this?”

Dex Hunter-Torricke: The technology itself should allow us to do unbelievably transformative things for our societies. There is a debate about whether access to the technologies that are coming in the next 10 to 15 years could allow us to

build a radically different economic model where you have an abundance of resources and you are able to bring a whole bunch of new tools to bear on solving some of the massive fundamental challenges that we have talked about: climate change, inequality, poverty, and transformation of public services. AI that is as good as human intelligence, potentially, on most cognitive tasks, could be enormously game changing. We do not, however, have good answers for how to organise that and ensure that the rewards and the benefits accrue to our entire society, and not just to a tiny sliver of companies and the leaders of those companies. That is not a topic of mainstream conversation and we are right to be concerned about that.

The positives, though, are things that could dramatically change the quality of life for future generations. AI is not just one thing on its own. It turbocharges discoveries and breakthroughs on any number of fronts. If you could cure cancer, if you could enable commercial nuclear fusion power and get clean unlimited energy—how transformative that would be for our societies and our entire economic model.

Kevin Stewart: I will return to your earlier important point about the one-person, billion-dollar company. We already have on the planet billion-dollar—trillion-dollar—companies that are at the forefront of all of this. Some would argue that they are not ethical now because they do not pay the taxes that some of us believe that they should. You talked earlier about curing cancer and the possibility of new treatments coming into play, and we can already see the huge differences in terms of early diagnosis by AI applications. You asked, however, who those treatments would be available to. Will they be available only to the elites who run the big companies or will they be available to everyone? Those are the questions that we need to answer in order to deal with the pessimism about where this may leave a lot of folk out there.

Dex Hunter-Torricke: I strongly agree with that. We do not solve our problems by pretending that they do not exist. There has been a pretty one-sided framing of the value of this technology from folks who, I think reasonably, in a well-meaning way, want to champion its potential. There are massive challenges, however, that speak right to the core of what kinds of people and societies we want to be in a future where the most powerful technology in history is being summoned. We do not have good answers for those things.

The future is arriving at a much faster pace than any technological transformation that we have seen in history. This is not the industrial revolution or the arrival of the internet; it is something much bigger and faster. We need to have that conversation now, because otherwise the

technology will decide for us. The pace at which we adapt to these technologies is too slow. Very frankly, it is much too slow in Europe and the UK. How many people think that we managed the arrival of social media well? This is much more difficult and bigger and so we need to move quite quickly otherwise we will absolutely end up in a model where it is a tiny sliver of the global elite who reap all sorts of rewards and the average person may find their quality of life plateauing or declining. Then you will see the political and societal backlash that ends up threatening the system in a bunch of different ways.

Kayla-Megan Burns: I mentioned that I am conscious that, currently, the decisions about what is good, what is safe and what is fair in these models—which are being used globally—are in the hands of a very few people. There is a severe concentration of power. This week, even the head of Anthropic, or the head of OpenAI, came out and said that they were incredibly uncomfortable because they realise that they are the decision makers behind all of these types of ethical decisions that are having global implications across vast populations.

10:45

Dex has just outlined a scenario in which either we let the tech go unchallenged—we let it do its thing—and we wind up with severe concentrations of power in the hands of very few people and controlled by companies within the global elite, or we regulate. We need to not be afraid of stepping in and getting things wrong, maybe slightly overregulating, then slightly underregulating, and then finding the middle. We really need to lean into that zone and not be afraid of doing that and not stall in the fear of getting it wrong.

There are fantastic opportunities here. Take rare diseases as an example. Dex mentioned curing cancer. That is definitely possible, but some cancers get a lot more funding than others because they are much more common. Rare cancers get very little funding and although they might be technically easier to solve from a medical perspective, from a treatment perspective, because of the few people who wind up with those cancers, that research is not being done. Such research requires millions of pounds, a lot of teams and a lot of equipment, whereas in the AI era it might require one very smart person and an AI model and that could literally cure that rare cancer in much less time, which would be absolutely amazing. That would be transformational for so many people on a personal level but would also have wider societal implications.

I think that we need to get in and make sure that AI is regulated appropriately so that we do not

wind up with unreasonable concentrations of power that will have negative impacts. Instead, we must make sure that AI is a positive thing for society overall.

Kevin Stewart: I think that you talked earlier about AI-induced chaos. I am playing devil's advocate here, because we have to in some regards. Earlier, you held up your smartphone; Dex also asked what we have done with that technology and whether it has been beneficial. I think that we may all agree that we have more communication, but is it meaningful communication? In terms of using AI, by means of which we can get more done, how do we ensure that what we do is meaningful? For example, probably every single one of us around this table is receiving a lot more mass communication that has been produced by AI. The temptation is, of course, to respond by using AI, which is not my bag, I have to say, at this time. I am not sure that some of that communication is as meaningful as it should be. We could see a situation where there is lots more communication, but would it be worthwhile, meaningful and make a difference to our society? How do we get around some of those things so that we do not get to the AI-induced chaos that you talked about?

Kayla-Megan Burns: I think that you are referring to my mention of AI-induced psychosis. That is a phenomenon with people who are engaging with AI for prolonged periods. AI hallucinates to a degree; that has decreased recently but, depending on what they are doing with the models, it can increase, it can fluctuate. The real danger with AI is that it is being programmed for engagement because that benefits the companies—it involves spending more tokens and so on. Companies profit from people engaging with AI.

The problem with engagement with AI, particularly for prolonged periods, is that it is generally programmed to keep you happy. Last night, I was speaking to some people from Anthropic who described AI as a very happy and affirmative entity. Whenever you let two AI models speak to each other they tend to escalate into nearly a Buddhist state of everything being amazing and just a series of wonders. That is very interesting to observe, but can be problematic, because we have smart devices sitting in our pockets allowing us to access all sorts of information, whether it is correct or not. It is out there and we can get it at any time.

People do doom scroll. That has had severe mental health implications and severe implications for democracy among all sorts of things. The problem with AI agents being programmed to be affirmative is that you can go to them with anything, including the wildest, most outlandish

beliefs, and although the agent might call you out—some of them have some safeguards in place and will call you out for unsafe behaviours—we find that the more you engage, the more you are able to work around those safeguards. That means that AI agents will affirm conspiracy theories, and dangerous or harmful beliefs and, in some situations, they will encourage dangerous behaviours. That is a real risk and a serious escalation from what we have on social media, because it is affirming and creating a one-person silo, which can affirm dangerous beliefs.

There is a serious risk of that, and that is again where regulation needs to come in and we really need to dive in. We need to deepen our understanding of how AI works, because currently we do not know how the models work. We do not understand their logic paths and their thinking. From speaking to a lot of the leaders in the industry, I think that there is generally a belief out there that we will have a lot of breakthroughs in neuroscience and psychology within the next few years, because we are spending so much time looking at how neural networks work, what the AI thought patterns are and how AI agents reach their decisions. We do not currently know that, but industry leaders genuinely believe that we will have a lot of breakthroughs in human neuroscience and human psychology because we will spend so much time looking at the models.

Kevin Stewart: We will wait and see whether we have those breakthroughs. Dex, could you also answer that question and maybe go a little bit further? We have had a discussion about the guardrails, the safeguards and what we need to do there but, as has been discussed, what is really required is an international framework agreement, which I think may be unlikely or very difficult to reach. How do we persuade the elites who are in control that an international framework is the right thing for all of us, and also for them, to follow?

Dex Hunter-Torricke: On the international framework, a lot of the things that I am sure that pretty much everyone in the room would want for a strong, successful future Scottish society are the same things that countries all over the world are also really hungry to hold on to. You want a future where you have economic growth and strong resilient societies. I am sure that you want a future where you have a well-equipped younger generation with the right skills to navigate the world and that is prepared for the huge range of challenges to come. We want to protect what it means to have Scottish culture and values—the whole set of things that makes Scotland special. Those are things where I think that there is a conversation that is ripe to be had with leaders from a lot of other countries, including a lot of smaller countries that also feel that they do not

have a seat at the top table for some of the existing AI governance conversations.

There is obviously some semblance of international engagement between the major powers over AI. There are the AI safety summits, one of which took place at Bletchley Park; an upcoming summit is being hosted by India. Those are established forums and they have a certain type of value, but a bunch of other folks have to have a seat at the table. There is a huge opportunity for Scotland to look at driving intensified collaboration among other countries that are in the same sort of position and also have those kinds of concerns—countries that might not be in the top rank of the powers in the existing AI governance frameworks—and to say, “These are all the things that we also want to see as part of that framework.”

There are a lot of those countries. If you add up all the powers that are in a similar position, they represent well more than half of the world's population. There is a big, open diplomatic opportunity right now. Over the past couple of months I have been travelling all over the world and talking to policy makers, industry leaders and civil society leaders. I have had a number of those conversations and I think that there is a great willingness to look for new models of international engagement entirely on these things. There is no perfected home or framework for talking about a bunch of them. A lot of the AI safety summit was literally focused on the technical safety risks of AI models. Everything that we have talked about is just scratching the surface of the societal challenges of AI; it is about vastly more than just safety.

Kevin Stewart: The convener is desperate for me to finish, because we are running out of time. I have one final question. We have talked about the Scotlands, the Estonias and the South Koreas of this world, but the drivers of all of this—the elites, if you like—are mainly American tech companies. America is not a nation that is renowned for creating good regulation, and it could be said of the current regime in the United States that that is even less so. How do we persuade those giants, or the elites, that an international framework is also the right way forward for them?

Dex Hunter-Torricke: Public polling in the United States on the attitudes of a lot of different leaders, including from public society, shows a great deal of commonality with a number of the things that we have talked about. There is a great deal of concern about the future that we are heading for. Big tech and the current US administration have a particular point of view on where they want to get to with the AI ecosystem. I think that it is a view that is not very popular with a lot of people and so there is a great opportunity—a

window of time—now to begin to drive this conversation in a much bigger way across different ecosystems. We will see how that can begin to shape some of the decision-maker preferences within the tech industry. Fundamentally, at the end of the day, the tech industry requires society in order to be able to fully function. The conversation on these things has yet to become mainstream, so we do not fully know where we can shape this.

The Convener: This has been a really incredible session. There have been a lot of very interesting answers. I am also intrigued to figure out which time zones you are in: the sun has been setting for Dex and has been rising for Kayla-Megan, and that has been interesting to watch. There are some very interesting things that we will definitely want to follow up on, so thank you very much for your time—I was about to say this morning, but this evening or this morning, whichever is applicable to you.

Dex Hunter-Torricke: Thank you.

Kayla-Megan Burns: Thank you for having us.

The Convener: I suspend the meeting for 10 minutes; I ask members to be back for 10 past 11, please.

10:58

Meeting suspended.

11:10

On resuming—

The Convener: Welcome back for our second panel of witnesses for our short inquiry into artificial intelligence. I am very pleased that we are joined by Steve Aitken, the founder of Intelligent Plant Ltd, and Leo Fakhrul, chief executive officer of XYNQ and Mamba Sounds. Unfortunately, Rich Wilson is not with us. I note on the record that he had a small accident this morning. I am sure that all members of the committee join me in wishing him a speedy recovery.

We have just over an hour, so I would appreciate concise questions and, if possible—although this is an expansive topic—concise responses. I will hand over straight away to my deputy convener, Michelle Thomson.

Michelle Thomson: Good morning. I thank both our witnesses for joining us. I will come to you first, Leo. Originally, our papers showed that Ziyad, who I think is a partner of yours, was to appear for Mamba Sounds, but I think that you are appearing under a different company name today. It would be useful, first of all, to understand what you are doing in the AI space and why, and what has brought you to this point.

Leo Fakhrul (XYNQ): Thank you very much. Yes, Ziyad Alrasbi was meant to appear in front of you today.

The problem that we are solving at XYNQ is that of fraud in the music industry. There is an almost \$3 billion worldwide problem at the first mile in the music industry. XYNQ is positioned directly to solve that problem by using artificial intelligence, as well as machine learning, to find the bad actors before the fraud even happens. Competitors in this space deploy fraud tactics once a fraud has already happened, and Mamba Sounds, which has already operated as a record label and an artists' collective, has data and infrastructure to stop the fraud before it even happens through XYNQ, the spin-off company. I hope that that makes sense.

Michelle Thomson: It makes complete sense. This session follows our earlier session with Kayla-Megan Burns, who is a board member specialising in AI for the Royal Scottish National Orchestra. It would be useful to understand the scale of the problem and the implications for the people in the artistic sector of fraudulent activity around their material.

We also heard from Dex Hunter-Torricke in our earlier session, who said that he could see the possibility of one person operating a company that would have turnover of \$1 billion with effective utilisation of AI.

It would be useful to understand the scale of the problem, where you see yourself operating and why you think that the new product that you are looking at could fit into that niche.

Leo Fakhrul: On the scale of the problem, we need to look at the global music industry first—specifically, the digital music industry. By 2030, the global industry for digital music specifically will be worth more than \$100 billion. At the current time, 10 per cent of the activity across the industry is fraudulent so, by 2030, the problem will be worth \$10 billion worldwide. You are looking at hotspots such as Indonesia, Brazil and South America overall, and parts of south-east Asia as the problematic high-ground areas for fraudulent activity.

What do I mean by music fraud in the first place? I am talking about metadata spoofing. An artist or a bad actor comes in and gives out bad information that is not official—it is not authentic. They use fake IDs to upload music under an alias. For example, artificial music might be generated using Jay-Z's voice for a fake song that never actually happened. The intellectual property should belong to Jay-Z himself, because it is his voice, but the AI has made the music, uploaded it and can make money from the song.

11:15

One aim of the fraudster is to launder money. Every single day, 120,000 songs are uploaded through distributors on to digital streaming platforms such as Spotify, Apple Music and YouTube Music. The number is looking to rise to 200,000 every single day by 2030. The scale of the problem is magnificent, because no distributor company has the time to look through 200,000 songs per day.

However, although AI can be used to commit fraud, we can also use it to retaliate against the fraud. We can use machine learning to find the anomalies within behavioural data and the metadata, as well as network economics, to be able to understand where the fraudsters come from and put a stop to the fraud at source. XYNQ is specifically designed to stop the problem and we have a speciality within the music industry. We have been operating as artists, a label and a distributor for the past three and a half years as the company Mamba Sounds. The spin-off has given us the direct tools, through Techscaler, The Data Lab and other Scottish organisations, to be able to position ourselves in the global market to tackle fraud in the music industry.

Michelle Thomson: Thank you. I will open this out to both of you, given that Steve Aitken has a very established company. I would like to finish off by exploring what you see as the critical factors in terms of skills and the ecosystem that have enabled you to operate as you do and which, critically, could enable Scotland to compete globally in this area. If we think of other industries, we cannot compete in certain areas at scale, but this is an area where we can compete. I ask Steve to answer that first. I have looked at your background, so I know what it is.

Steve Aitken (Intelligent Plant): You asked about skills and the ecosystem, and how that has meant that we have been able to have a computing science company that has lasted for just about 20 years. I went to university in Aberdeen to study philosophy to begin with, but changed to computing science in the end because I figured that it would get me a job. There is a slight irony with what is going on now when you think about it. However, without that ecosystem and the knowledge to get on to the first step of work, you can struggle to get started.

From there, I worked for an employer for what felt like a lifetime, although it turned out that it was only four years before I set up the company. You can get a lot from that experience: you get the exposure to the economy and to what companies do and how they operate and keep themselves going, which is important.

That kind of experience is close to my heart. We run the inform prize, which is about getting students and the projects that they are doing in front of companies as early as possible so that the companies can see them. We will be bringing the prize to the rest of Scotland next year. We were in Aberdeen and Glasgow this year; Edinburgh is well up for it; and we also have Stirling and a number of others. We have gone for every university that does computing science.

On the ecosystem, I think that Scotland is good at people coming together and being able to do things together. When I say "people", I mean people and companies. The only way that you can do that is with something that is key to what the committee is discussing, and that is trust. Trust is at the root of all the ethics. Why do you have ethics? Because without ethics, you cannot have trust. We are good at building trust with people. I think that Scottish people are good at that because we can be straight talking and we tend not to overegg things, which is good for building trust. That has meant that I am surrounded by a number of people and a number of companies that I honestly believe are out for all our best interests.

The skills to be able to do that are things that I have learned rather than things that I have been taught. Ethics and how to behave correctly are becoming very important in this field, and bringing ethics to the right levels would make a big difference.

I remember being in front of a business school class when a student put up his hand and asked, "If I have the choice between doing the right thing and making no money, or doing the wrong thing and making a lot of money for the company, what should I do?" I said, "If you do the wrong thing, you might make a lot of money today, but you will make nothing in the future." There was a lot of surprise about that comment to a business school class. I would like to bring in that surprise a lot earlier in people's learning so that it is really grounded in their character.

Michelle Thomson: You make an important point. Before I was elected, I did some primary research into the perception of Scotland's global diaspora, with about 1,200 participants across 72 countries. One of the big themes that came out was about the trust factors in relation to Scotland as a place to do business and Scots as people to do business with. That is something that we can trade on, because it is a currency that has high value in today's world.

I return to my point. Leo, you are younger and we all hope that you have a great career and future ahead of you. What have you seen in the skills that you have been able to learn and the ecosystem that has supported you that gives you

confidence that Scotland can compete globally in this space?

Leo Fakhru: My co-founder and I began at school, where we had the opportunity to do national 5s, highers and advanced highers and then make our way to either college or university. I went to college before I went to university. My background is in economics and finance. I was an undergraduate for free, without any student debt, and now I can go back to university through The Data Lab's scholarship incentive and study artificial intelligence and data science, which is a huge boost not only for my skill set but for the company. That kind of foundation is not found everywhere. I am super grateful for everything that Scotland has been able to do for me and my family. I am also grateful that we can come here, build something so great, compete at a larger scale against multinationals from Canada, the US and all round the world, and have a massive say.

You just spoke about trust with Steve, which is huge. In the past three and half years, my co-founder and I have built deep trust with the industry—with distributors and our clientele—and that explains how we have an organic pipeline that is ready to go as soon as the product is already. We already have letters of interest at hand to be able to make this a real product with a unique value proposition. It is one of the things that keeps us in front of our competitors.

I would urge a lot of youngsters in Scotland to use the infrastructure and ecosystem that we have in front of us. Go to university and grab a degree in engineering, medicine or economics—even law—if you want to study, but use that degree in any way, shape or form you can in the future, and try to use it for the good. For us, that has been massively helpful.

My co-founder has a computing science background and has been working in London. I wish that he could work somewhere in Scotland and could compete directly with our London competitors down the road. I think that the ecosystem is a massive testament to what this country has been able to do with companies such as Skyscanner that have come out of the ground. I want XYNQ to be one of those companies alongside Skyscanner.

Gordon MacDonald: Good morning. I will ask about governance and regulation. It blew me away when Leo Fakhru highlighted that there are 120,000 fraudulent music releases every single day. What is the Government's role in addressing that? Given that this is happening across not only countries but continents, what should be the Government's role in helping to stop the fraud that you are talking about?

Leo Fakhru: I make the point that it is 120,000 songs that are uploaded every single day, and about 10 per cent of them are fraudulent. However, given the current rate of developments, we will probably soon be looking at 20,000 fraudulent songs being uploaded every single day.

The role of governance will be massive. We are looking at digital streaming acts across Europe, as well as in the US, which are coming into hand quite massively. What we are looking at for the product that we are building is based on those rules and regulations. When the EU's Digital Services Act comes in, we are expecting all the companies across Spotify and all the digital streaming platforms to use those regulations directly to stop fraud at the first mile, which is where XYNQ comes in with the product Redflag. That is our strong point and our unique selling point. The role of governance is massive.

I was reading about how Singapore plays with AI in its ecosystem. People there are not building the next ChatGPT, but they are building the regulations around the next ChatGPT. They are putting limits on what is possible with the likes of these superpowers that we can work with today. They are also taking the view that, if they have trust, security and safety in the ecosystem, they will be one of the leaders in the system, although they will not be like ChatGPT or OpenAI—they will have a whole different angle.

There is a requirement to build such a system. In the financial sector and in banking—wherever you are—fraud is a massive concept. In Scotland, we have the infrastructure to improve safety in those markets, and I would love that to apply to artificial intelligence and machine learning, too.

Gordon MacDonald: The public have a lot of concerns about data privacy and data harvesting in relation to AI—we are talking about bias, misinformation and so on. Is there any legislation in place to tackle that anywhere in the world, which we can learn from if we are considering passing legislation in Scotland?

Leo Fakhru: We have to go back to Singapore, which is using a good tactic to protect its residents and citizens. On data privacy, the world wide web is a whole different world that connects everybody a lot more quickly. We need to look at positioning Scotland with a voice to tackle the issues directly and have an influence at the table when such conversations are happening. I would love to know Steve Aitken's thoughts on that.

Steve Aitken: I am in a place where I currently have a thing going against a company under the general data protection regulation. It is key to any legislation that it is able to be enforced, because it is far too easy to have something that is quite wide

but does not actually have teeth. I am wondering which way my situation will go.

Gordon MacDonald: What should the focus be on?

Steve Aitken: It should be on what you can actually do. If we put in place regulations, they should be on something that you can verify and which you can prevent from happening or make happen.

In this space, AI and large language models have what is called a system prompt—that is the bit in the background that is given to AI before it is used, which tells it how to behave, what to highlight and what to hide. The system prompt often says, “Be ethical and do these things in the right way,” and that sort of thing. However, not everybody publishes the system prompts that are being used, which are tilting everything that is being responded to. In my mind, those prompts should be published.

Gordon MacDonald: I think that I will leave it at that.

The Convener: Terrific. I would like to bring in Willie Coffey.

Willie Coffey: Good morning. I will start with Leo Fakhru. We are always going to need XYNQ because, by the sounds of it, we are always going to need to retaliate against the bad-faith actors. The committee was talking earlier about whether to embed the ethical approach and whether that is possible. I will come to Steve Aitken in a moment to ask more about that.

Leo, without giving any of your secrets away, can you say whether we can successfully do what you are setting out to do? Will we be able to prevent fraud today, although it will reappear in another form tomorrow? Will it be an endless journey for companies such as yours to retaliate against fraud? Is that what we will be seeing now and into the future—a constant fight between good-faith actors and bad-faith actors?

11:30

Leo Fakhru: Yes—that is a legacy problem. We have had that since the dawn of time in whichever industry. What I see for myself and for our company is that artificial intelligence is making music that is passing the Turing test with flying colours. The Turing test concerns how well the AI can pass at displaying human behaviour—I believe that that is a simple way to put it. Such music is passing with flying colours. That means that AI music sounds just as human as human-produced music, which is a huge problem.

How do you detect AI music? That is the question. For us, that is about how we evolve the

product as time goes by. Evolving the product is directly focused on a few things, such as the metadata and the network in which the music is uploaded.

Without giving too much away, we want to input a know-your-customer system, but we want to call it the know-your-artist system, where we identify each and every person who is uploading their music and then see who they are connected with the network economic effect. On Spotify, we can see who somebody has collaborations with. If I made a song today with Steve Aitken, I would put Steve as a collaborator, as a feature. If I committed some level of streaming fraud—for example, by paying to inflate my numbers—Steve would also be impacted and would receive a negative score on his profile.

Where Redflag at XYNQ comes in is that we would look at what Steve Aitken had uploaded in the previous few months. What was his pattern of behaviour? What devices was he uploading from? Where was the music coming from? Were we recognising his voice every time, or was the voice AI generated, so that we were listening to a bunch of AI slush, as we call it? That is where we want to come in. We want the product to evolve as new AIs bring new features towards us, but it gives us a good, strong seat to tackle the problem.

The data raises a massive issue. First and foremost, we need to get our hands on the data. We do not have enough distributors in Scotland, which is why, when we initiated Mamba Sounds, we focused on a worldwide approach. We focused on music in Africa being supported in the US, the UK, Germany, France, Holland and Canada. Those are our top seven markets. From distributors that are working directly in the continent of Africa, we took music from Africa all the way to a global market in the west. That gives us the advantage. We win against everybody else because we have direct access to the global distributors before anybody else does. I hope that that answered your question.

Willie Coffey: It really did. What is the public's attitude to slush? If it is cheaper, do people care?

Leo Fakhru: From a consumer perspective, I sit in both seats—actually, I sit in all three seats, as I listen to music regularly, because we get 10 to 20 submissions a day worldwide for Mamba Sounds. We also look at the situation from other seats. We sit in the fraud and regulatory seat, but we also consider the consumer aspect.

There is a place for AI music, but the difference is that it is not worth paying for. On YouTube, you could look up a song in the style of Frank Sinatra but maybe in a different version—a rap version. There is a time and a place for that, but it should be in a different market from Spotify, where you

have to pay to listen. A lot of digital service providers are now looking at completely banning AI music, but the question at the beginning is: how do we differentiate between the two? As I said, AI music passes the Turing test with flying colours. Providers want to ban AI music, but they do not know what AI music sounds like in the first place.

For Scotland and for XYNQ, we are looking at methods for harnessing a fair playground in which AI music can come through, but subscribers who pay every month do not have to pay for it. If that was the approach, there would be a whole different market, which could be similar to Spotify. Maybe something like Almusic.com would come along—this is completely fictitious—where people could listen to AI music or generate their own AI music.

I will give you a real reference point. There is an AI company called Suno that allows people to type in a prompt to develop their own song; people can harness their own beats using AI and give their song styles and so on. I know many big producers around the world who are taking the stems from that product and rearranging them to produce a sound that they enjoy. Is that artificial intelligence music or is it those producers' music? That is a debate to be had at another time.

Willie Coffey: What about creativity?

Leo Fakhrul: To talk about another aspect, when Auto-Tune came out in the 1990s, people said that it would ruin music completely, but artists such as Kanye West and T-Pain have been able to use Auto-Tune—I am thinking about R & B and hip-hop music specifically. They used Auto-Tune in a way that was quite artistically sound.

AI music is now coming through the ranks, but artists and producers can use it directly in their own way. What I am very against is the AI making and uploading the full song and receiving royalties from it. That raises a major problem for artists. On average, an artist on Spotify will have received \$12 of income per month in the past year. If that pool of income was challenged by AI, the amount per person could become very limited. In the past year, one AI account was able to make \$10 million—I repeat: \$10 million—from uploading AI jazz music. Let us think about how much money was taken away from the average artist, including people who are actually making jazz music, by that one account.

I know a lot of orchestras across Scotland that are producing music. They are perhaps not uploading it to Spotify but, if they did, they would get a lot of welcome support. However, they are competing against AI now, which can artificially make a song in seconds, upload it and frantically money launder. The amount of money that can be made from uploading the result of a few seconds

of work is stupendous. That undermines a lot of the artistic values that we hold, especially in the creative sector in Scotland. This is a massive problem, because we are a country that is proud of what we can achieve and what we do. We have the likes of Lewis Capaldi and Calvin Harris—guys who are at the top of the game. If AI music came from them, would we be happy with it? What do we count as AI music now? That is the question to you from my end.

Willie Coffey: Thank you for that, Leo. Turning to you, Steve Aitken, and the wider issue of the ethical battle, should there be an ethical blanket thrown over the whole AI revolution? Is it possible to do that? Is it always going to be the fight that Leo Fakhrul describes? Can we win that battle, and should we try to win that battle?

Steve Aitken: Ethics is always a fight between what is right and what is easy, and that will always be a thing. The key thing is people rather than the technology. When we think about the ethics, it is making sure that people are doing the right thing with AI. If you consider that it is a massive multiplier, that is probably where the issue is: if people are doing the right thing, they will be doing it better with AI; if they are doing the wrong thing, unfortunately they will also be doing it better with AI. I think that we need to focus on people and make sure that when they are doing the wrong thing we hold them to account and that when they are doing the right thing we reward them.

Willie Coffey: That goes into the area that Gordon MacDonald led us into—regulation, control and standards. Is it too late to try to establish that stuff?

Steve Aitken: In some way, that will already be established in law and the way that we deal with directors of companies. It is about understanding where the power lies and who it is that can make the most use of it, and making sure that there are checks and balances in place. I think that there could be more in that space, but that would be applicable with or without AI. It is just that the risk and opportunity are both higher with AI. It might make us focus a bit more on making sure that people who run companies run them ethically.

Willie Coffey: You may have heard Dex Hunter-Torricke telling us earlier about the advent of the corporate billion-dollar company with a single person in control. How do we persuade such a person to embrace an ethical framework and ethical standards? Is that a journey that we just have to keep working on and fighting to achieve?

Steve Aitken: The single person most people would think about in that context would be Elon Musk. If you look at his rise, most of it was through at least a facade of ethics. He looked like he was

doing the right thing, and he said that he was going to do the right thing. Because of that, he got a lot of backing. I would say that since then he has taken a slightly different view—I have a sticker on my car that says, “I bought this before he went nuts”. People buy things from people who they like and trust, and when that is broken, then it falls down. If someone wants to build a billion-dollar company and then they are done, great, but what will they do next? We need to make sure that people can see beyond that.

This is another thing that is close to my heart. The first company I worked for—for those four years—had two directors, who sold the company a few years later. After selling the company, the directors had a big existential question: what am I here for? What am I supposed to do? I have all this money but what will I do with it? Did I prefer to have the company rather than the money?

We need to bring such realisations to people at an early age so that they can see that building a company is not just about earning a lot of money, but is about doing a good thing. When you come to the end of your life, everybody cares about how they are viewed and whether they are seen to have been a good person. There is no amount of money that can buy that.

Willie Coffey: Trust is at the heart of this. Is trust our saviour? Is trust going to save us from a horrible future where we will be endlessly fighting against bad-faith actors? If we can establish that within any of these frameworks do we have a chance, Leo Fakhru?

Leo Fakhru: When we first came up with the concept of the fraud problem itself, it was because we were defrauded in the first place. Somebody went ahead and put Mamba Sounds as an artist collaborator on a song that we did not make or sign off on. We were hurt financially and reputationally as well as operationally. It affected what music we put out later on.

Going back to your question, the reason that we do these things and make these companies is ultimately to answer or solve a problem. I think that most entrepreneurs, especially the people who I hang around and spend time with, are looking to solve a problem and make lives easier. A good friend of mine owns a company called ScrubMarine, which is a maritime technology company. He says that they are looking to give a painkiller, not a multivitamin—it is meant to stop the pain for the industry, not just be a feature for a product to make you feel okay, like a multivitamin.

When we speak about these solutions, we are all using AI in some way, in our day to day and in our solutions. It comes down to the problem that you are solving and why you are solving it. That is one of the bigger asks for a lot of entrepreneurs

and people who can raise funding. When you are sitting in front of an investor committee and building those relationships, they will ask you a few questions—“Why are you building this in the first place?”, “What is the problem that you are solving and how are you solving it?”, and so on—and you want to have some level of resonating feeling in respect of the problem. When we solve a problem, we are solving it from a place of hurt—we do not want it to happen again. Imagine if we had a system such as Redflag that stopped the fraud before it even happened. We would be safe.

For me and a lot of entrepreneurs who are solving big problems within Scotland just now and in the UK overall, we have to look at how we are impacting society. I had a problem with it at the beginning. In the past few weeks, I was asking myself who our product is actually benefiting. Is it the distributors? Is it the big labels? Is it the Warners, or the Universal Musics? I have thought about it and even the guys who we are helping right now—the independent artists—have a level of satisfaction for the product. We started interviewing a few artists and so on, and we said, “What if we were able to give you a verification—a pass, an integrity badge, or a licence—that says that your song is fraud free, which you could take to Spotify?” Let us say that Spotify had false allegations against you. You could use that verification under the law or in the court system, because our data would have a 0.001 per cent false positive rate. We want it to be at that high level so that we can say, “This person’s song is fully fraud free” and we are hyperfocused on making sure that it is real artistic, creative content.

Our belief that we can go back and help the small, independent artists sparked something within me that said that we are doing the right thing. We are helping the ants in the system who help to build the foundation as well as the big players—the lions, the kings of the jungle, essentially. Everyone can benefit from this. When you are solving a problem, you have to look at who is benefiting. Who are you aiding? Who are you stopping the pain for? For us, it is industry-wide. It is on that vertical: artists, distributors, labels and digital service providers, such as Spotify and so on.

11:45

Willie Coffey: Okay. Absolutely fascinating, guys. Thank you.

The Convener: Michelle Thomson would like to follow up on that.

Michelle Thomson: I want to follow up a point that Steve Aitken made. Ethics has been a golden thread running through the three evidence sessions we have had. We have heard the need

for that emphasised by a variety of witnesses. If we forecast forward, AI is potentially a significant disrupter to our society. For many people, having more leisure time is a curse as well as an opportunity. On your point, therefore, about ethics, does it mean that from a skills perspective that we should be teaching ethics in schools because humanity—and I appreciate that this is quite a big question—will have to encounter this existential crisis, arguably triggered by significant momentum in AI? In this committee we are not going to solve any of that, but should we be thinking practically about teaching more ethics in schools to counter some of this?

Steve Aitken: You are saying exactly the stuff that I am thinking. When you look at what is happening, you see that that existential question that the director who sold the company had is coming to us all. The end result of loads of leisure time might mean realising that dream of being able to have coffee with the sea in front of you, but then you go, “Why am I here? What am I doing?”

The arts have taken a back seat to science in a big way in the past 10 years, but you can see a resurgence in the arts being able to bring that to people. I think that learning about ethics at an early age would make a big difference. It is not about arts or science on their own, it is about how they combine and the way you can say, “Here is how you can do the right thing, this is where you will gain, this is how other people will gain, and this is how we will all live”. I think that we need to get that at an early age. It might come down to even how we advise parents at that stage, because there are some grounding concepts that are important early in a child’s life.

Ultimately, we need to consider how we make sure that people can live. It is easy to look at the big questions around energy and the cost of heating and things like that. Those things have been right at the forefront of most people’s minds recently. When people struggle, they lose sight of the longer term. In that situation, the only thing that you care about is your survival, and you will make decisions that you will potentially regret later. We need to help people to see into the long term so that they can make the right decisions all the time.

We have a lot of people in the arts in universities who think that AI will replace them. I was chatting to some of them at Aberdeen university and I was saying that I think the complete opposite: I think that it creates a massive demand for the arts. I started off doing philosophy and went into computing, so I just see it going back the other way.

I strongly encourage bringing that education in earlier. It should not replace STEM. We need to get more STEM in for younger children, but we

need to add arts and that level of thinking to it so that it is not without a purpose.

Kevin Stewart: Perhaps a computing science degree with philosophy is the way forward for all this.

Thanks for coming today. Convener, I should say that I have met with Steve previously for some good conversation.

I will stick with the ethics aspect. We have heard from others today about who is doing well. Leo Fakhrul, you mentioned the fact that Singapore is driving things forward but with limits. We are operating in a global context here. What we require for governance is an international framework, which may not be seen as beneficial to some of the elites out there. Is such a framework required for us to have the right governance here and to continue the trust that there obviously is in Scotland’s businesses?

Leo Fakhrul: When I think of this atmosphere of artificial intelligence, the first thing to say is that it is rapid—it is fast. On any regulations and laws that you want to put in, you will have to move very quickly, otherwise you will be at the backhand. Secondly, we cannot tell other countries what to do. It is hard for us to make sure that they believe in what we do, because they might see things from a different standpoint.

However, we can make sure of regulations that are within Scotland. If we can do that effectively, I believe that we could have a playground—a sandbox, if you will—for a lot of the high-growth companies to come to Scotland and build here in an ethical but limited way. Those are the voices that you want here, who can help to shape worldwide regulations. If you can bring more companies here—if you are able to put the policies in place to bring the leaders in this space to build in Scotland—you will soon realise that Scotland has a massive voice at the table.

We talked about how one individual could run a billion-dollar company. How do you bring that person to operate in Scotland? To harvest that framework, and that network, it is important to say that we have the facilities to be able to build this. That is an infrastructure thing that we can do from here on out. If we can build in Scotland the data set, data rooms and so on that we need, that will give Scotland a massive advantage worldwide. We can attract more talent to Scotland and build from within. When you build from within, you will have a taller tower compared with those of your competitors.

Steve Aitken: I was just going through in my head the way that this tends to work with investors and founders and how much control a founder really has when there is an investor. I have not seen a lot of thought put into the ethics of

investors. I have seen founders who are full of ethics being pulled in directions that they do not want to be pulled in by their investors. I do not know what the answer to that is.

On whether we need a global framework, it would be really helpful if we had it, but I worry that it might not come in time to have the effect that we need. The global framework is the stick, but the carrot would be easier, can be done a lot quicker and can be done locally. By taking the carrot approach of helping people realise how ethics work, and by helping them set up companies that embody trust and that will thrive in an environment where people will choose them, perhaps we are the country that will have the billion-dollar, single-person company that then decides that it wants to make the rest of the country a good thing.

Kevin Stewart: I do not want to put words in your mouth, but are you saying that, in that carrot situation, the decisions that the likes of Scottish Enterprise or the Scottish National Investment Bank make about helping out new players in the game should have an ethical basis built in before any investment comes into play? Is that what you are suggesting?

Steve Aitken: That would be brilliant. At the moment, a lot of those decisions are made in the same way that an investor would make decisions: based on concern for whether the money will create more money as opposed to whether the money is going in the right direction. That is my opinion. I think that it is because we set the rules that way and people follow them. We say to them, "This is public money. Treat it like your own. We do not want you to waste it," so people do their best to say, "How do we make sure that we keep this?" and they tend to look, in my view, at the short term: "If we invest this, do we get this back in the next two or three years? How does that work?" That is a common approach: it is not particular to any organisation. Only by explaining example stories can we change those opinions.

However, adding that as a thing to think about and making the instruction to those organisations could have a massive impact. Being able to say that we are looking to encourage businesses that are doing the right thing, that are helping others, and that are doing things that the country would be proud of would have a massive, positive economic benefit.

Kevin Stewart: Would you agree with that, Leo?

Leo Fakhrul: A hundred per cent. That was very well put together, Steve. What you said, Kevin, about needing the Scottish National Investment Bank to make those decisions from the top down, is also a good strategy. What investors are we seeing? What angel investors are we

seeing in Scotland? Where are they coming in from? How are they holding themselves and respecting themselves in this industry? I think that that is where it begins, as well as having really good relations with investors. Prior to Steve's answer, I was going to say, "Find investors who are on the same wavelength as you," but that is not always possible. To find investors who are on the same wavelength, want to make money and have the right ethics is almost impossible.

I know a lot of founders in Scotland, specifically through the Royal Bank of Scotland accelerator programme as well as Techscaler—we were also part of Heriot-Watt University's business school incubator—and they are always looking for investors. Founders are always looking to the US because the ticket sizes are huge. We can look at the US and they will come in with huge cheques for us, but it might make us question what our ethics are and how strongly we believe in them.

XYNQ would benefit massively if the Scottish National Investment Bank were to look at us and say, "The problem you are solving is huge and we can back you. We can give you the right infrastructure and the right nest to build it in, but we want you to stay here." That is a fantastic offer that I would never say no to. You would have to look at that and think that that is what will keep us here and, at the same time, help us to develop something bigger in a much more strategic way. That is what we need more than anything. We need that strategy from within.

Kevin Stewart: Are any other carrots required to ensure that business that is being carried out here is ethical?

Steve Aitken: We spoke about education for young people. There might be the same thing around education for people who are starting up businesses and people who are investing in businesses. If we do research into how it plays out when you act in different ways and what the outcome is, that could give definitive answers so that people can scientifically say, "This is the right thing to do," not just because it is the right thing to do but because it is actually good for us.

A lot of it is about education, and I think that is true also of the use of AI and some of the other dangers around it. Our challenge with that is that education is changing rapidly. What we would tell someone this year has probably moved on from what we would have told them last year about what to look out for when you are using AI, things to think about and things to worry about. We need a very dynamic way to keep closing the loop of information from people who have knowledge in this space back to everyone—the public.

12:00

The thing is that everyone in the public has access to the likes of ChatGPT and Gemini, and I worry about that. I will give you an example. I needed to get three quotes to do a purchase. It is similar to some of the processes in the Scottish Government. I thought that I would get AI help with it. I typed a computer prompt: "Where can I get three quotes for this?" It came back and said, "I found these three places you could get it from. Do you want me to generate the quotes?" Knowing how AI works, I know that any quotes it generated would not have been generated by the companies; they would have been generated from nothing. Joe Public would accidentally be committing fraud. They would go, "Yes, please. Great, I have these quotes. They have the logos of companies on them. They have all the right numbers" and go ahead and use it. We really need to close the loop on those things by using education, with the right stories, so that when people get a result they know what it means, what is behind it, and what it is and is not capable of.

Kevin Stewart: Do you have anything to add to that, Leo?

Leo Fakhrul: On the subject of carrots and what will be a good incentive for us at XYNQ to remain ethical and so on, I look at it in terms of the benefits of staying at that ethical level. How can we go ahead understanding our decisions? For myself, it is by developing an advisory board that is not only complementary but full of diverse voices and so on, to help me make the best decision as CEO of the company. That gives us a good ground game, if you will, against the companies that we are competing against.

If you look at the carrot: yes, education is already on the table. The Data Lab is offering scholarships to universities and partnering with universities to give them access to the resources that they need to have that high-level certification. It would even also look at how we get talent who have access to this network, how we can get them into these cycles, and these jobs. We spoke about a one-man, billion-dollar company. Who wants that? Realistically, who actually wants that? People thrive in a network of people. We are nothing without one another, and I think that community is a massive thing. People work for the sole purpose of being next to people, I think. It is one of the things that we need.

How can we get the workforce in Scotland to work with each other, and how can we trade within each other as well? How can we support each other even more? Those are the benefits that I am looking at for XYNQ. How can we offset some of the responsibilities that we have to other Scottish companies? That would be massive. That would

help us with tax rates as well. Value added tax could be reduced a little bit in there.

Those are the things that we are looking at. If we can build it and develop it in-house in Scotland, can we get it cheaper compared with making it in India or even Turkey?

Kevin Stewart: Steve, you mentioned procurement and Government procurement. That may be another carrot. If we could look at Government and public sector procurement when it comes to AI and build in ethical standards there, would that be beneficial?

Steve Aitken: I think so. A lot of that will change going forward because the way that procurement is done tends to rely on people generating documents, which they will not be doing themselves any more. When you create the criteria and put that out to a number of people who have access to an AI tool, they can go, "There are the criteria. Generate the document." We have killed the market for people who were consultants on how to create a bid. I feel like that was not work that was well placed anyway.

The good side of it is that it means that there is a level playing field. Everybody can put in and get out exactly what they are doing. In a lot of places, people would be doing amazing things but would be so lost in the detail that they would struggle to get the right points into a bid. So, there is a good side to this: everyone will be able to generate incredibly good bid documents.

On the other side, I imagine that it will probably be something like that AI tool that is doing the analysis in the first place, too. Then it will come down to a human to make the decision. One of the things that runs through my head quite a lot is that, usually, the point at which we have a human in the loop is legally required. Computers do not make decisions. People have to, because then there is someone who is legally responsible for the decision. If a computer makes a decision, you cannot put it in jail.

Kevin Stewart: I will stick with trust but change tack a little bit. Obviously, there are huge opportunities for AI. A lot of that opportunity has not been grasped yet. A lot of folk do not trust AI completely. I will use an example that you gave, Steve Aitken, without breaking any commercial confidentiality. In our discussions, you talked about finding an AI solution for a company, but the company still stuck with the original project that it had in place because there was an edginess about the entire scenario. How do we ensure that we get trust across the piece to ensure that we get the absolute best out of these new technologies?

Steve Aitken: That is a good question and something that I wish I had known at the time that I was talking to that company. It was a few years

ago and at a much earlier stage than AI; it was neural networks. Thanks for the example. It took me a while to click which one you were talking about.

In that case, we were showing engineers a way in which they could get more out of an asset—it was an oil and gas asset; a lot of our client base is in oil and gas and, I hope, transitioning. It showed them that they could get a lot more out of the asset that they had. The engineers reacted by trying to follow the worst advice that it gave. It was really quite interesting to watch their behaviour. They wanted to push the system out, because they were worried about what it would do to their jobs. That is the natural human reaction that will be echoed throughout the country right now. That is not the right decision because, in that moment, when someone makes such a decision, they are limiting themselves by not doing it. If an individual or a company does not adopt AI, they are behind everyone that has adopted it. That means that small companies, individuals and big companies do not have a choice. It is a case of, “Do this or do not survive.”

In the long term, we have the big question: when we follow this path, as we must do in the short term, where do we end up? It leaves us with a big question as to how we make sure that we will catch everyone at the other end. Now is the right time to think about that, because we probably have four or five years to think about it. It might be that we have around a parliamentary term to think about it. That is the time period. It is really short, and we need to think about the other end.

Kevin Stewart: Again, it is about ethics and what we do as politicians, and as leaders, to fully utilise the technology but also to find other jobs for folk whose jobs may be superseded by the technology. We heard earlier that we could end up with fully automated industries, so we have decisions to make about what meaningful work we find for folk who are currently in those industries.

Steve Aitken: When finding work for people, it is about giving them value and making them feel valued so that they know that they are providing a contribution of worth. At the moment, that is mostly done through work. Before the industrial revolution, work would have involved lifting and moving, and we could not conceive that someone would be sat in front of a display on a computer, hitting keys, and that that would be work, because it would look very easy. We do not yet know what the next type of work will be. Ethics comes to mind because it is something that is very hard for the computer to do. However, we do not know what it is and we need to focus on that and figure out what it is that people will be doing to better themselves and others going forward.

The Convener: I might just ask a couple of brief questions following on from that—I am not sure that they are brief, actually. This morning, we have ended up hyperfocusing on single-person billion-dollar companies and they will probably exist, but it is probably something that looks a little less dramatic that will be more pervasive. Listening to both the earlier session and this one, I wonder whether the point is that we need to focus on ensuring that people create rather than just consume, and produce rather than just process. It is very difficult to predict precisely, but following on from what you have just said, are those some of the shifts that we need to think about to ensure that we are leveraging AI? If we are just consuming and processing, AI can do that much better, but we are not really going to be part of the value chain if we are just consuming. Is that the right way of thinking about it? Is that a reasonable conclusion to draw from some of the things that we have been talking about this morning?

Steve Aitken: It is not only a conclusion; it is why AI companies are getting a lot of investment. A lot of what people are putting into questions is well thought about—a bit like the questions here. You can see genius in what they are doing and then you realise something that you did not before. At the moment, most AI models have been trained on all the data that you can see on the internet—the lot—and they have no more data to train on, so they need more. They are ravenous for data. People are giving it away with their questions and with their, “Are you sure about that? What about this?” Yes, we should encourage people to do that, and we should look at how we make sure that, as a country, we are getting the most value from that.

It may be that we see it in the people and they are enabled by this to do more. I think that that is probably more the way forward than in the past, which has been more about trying to limit the information and saying that knowing the information when other people do not is your USP. I think that we are getting down to more core things that are your USP than just your access to those things.

The Convener: To help us steer through that, I have a few questions. Both of you engage with technology and run technology-based companies. As you think about AI and what that means for how you go about your day-to-day work, how does it shape how you organise around a problem, how you organise your businesses and how you seek to arrange things to meet your customers’ and your clients’ demands? I think that we all probably have a very 20th century model in our heads about what a business looks like—there is a chief executive and he or she has four to six vice-presidents or whatever the latest job title is, and they all have a column of people who report to

them, and there might be some horizontals. To my mind, that goes, because that is throwing people at an information problem. When you think about your business, how should we be thinking about organising and organisations? What should those principles be if AI is in the mix?

Steve Aitken: Our organisation is quite flat. That gives us a lot of advantages and a lot of disadvantages at the same time. In the past, we did a lot of business with oil companies, and they would send us a contract that was maybe 500 pages long. As a company that, at that time, included three people, that would have taken us quite a long time to digest and understand, and we could not really afford anyone who was qualified to look at it. What did we do back in the old days? We just signed it and moved on. Now people can use these things to help them to understand what they could not have before; it would not have been possible. Now, in a start-up that does not have an in-house lawyer, people have something that is almost competent and able to help them. It is the same with accountants. It comes down then to what the company is doing. If it is a company that has lawyers, now all the lawyers are doing less work, but more is getting done.

It is an interesting question. I think that it is one that we cannot change. It will just happen. It is about how we make sure that we steer the public through it in the most structured way possible, because it will change an awful lot. A lot of people say that this is a bubble. In terms of the investment in AI companies I would agree with them. When it comes to the likes of OpenAI and others, you can download the things that they are building for free and you can run them, so the value is not there; the value is somewhere else. If we knew the answer to that, it would be good to make some investments in the stock market. We know that it is a bubble, but that does not mean that AI is not massive and it does not mean that it will not change society. It really is.

12:15

The Convener: Likewise, Leo Fakhru, if you were to explain a 21st century company that embeds AI, what would you say were its organising principles? It is not about functional silos based around information, because the AI will do that for you. What do you think the organising principles are for an AI-based company?

Leo Fakhru: As Steve Aitken mentioned, we are at an early stage; the company was set up on 5 November, almost a fortnight ago. However, the ambition is to have a very flat structure, as Steve mentioned. Many of the great tech companies, software development companies and business-to-business software-as-a-service companies that

we look at are focused on a flat structure as well. For example, at Apple, there might be a vice-president of AirPods in Asia. We want to employ the same structure in XYNQ, with different products on the market, to make sure that we are solving the problems that are at hand and that we want to solve, but making sure that we have a specialist in every sector as well.

I predict that we will see fewer lawyers at legal firms such as Dentons and many more in-house counsel for a lot of companies. We want to have a specialist on board and they can use the AI tools to have their own interns, if you will, or associates. They can use these guys to have their own legal firm within the company. I imagine that that is something that we would do at XYNQ as well. We would, over time, have specialists in music law and technology law, and have them employ AI and make sure that we understand how to use it.

There is a good way to use AI. I spend a bit of time speaking to people my dad's age, telling them that there is a good way to use AI and there is a bad way to use AI, in terms of what output you want and how prompting works. I think that Steve will be able to tell you a bit more about that, but prompting is important when you are getting a decision from AI and in how you use it. I am sure that everybody here uses AI. It is in that structure that we want to have specialists connected to the AI within the system.

Yes, we will be using a lot of AI. However, we will have specialists sitting on top of those projects, owning and having responsibility and accountability for the problems that they are solving and what they are saying in the meetings with us. It is about having an artificial assistant more than anything else and not about replacing those jobs.

The Convener: I have one final question. Steve Aitken, you said something quite interesting about trust. The reason why it is particularly interesting to me is that it aligned with something that I encountered recently. It also relates to things that Leo Fakhru was saying.

I was in Singapore, and I met Enterprise Singapore there. It is always interesting, and I like looking at different Governments and their policies and their agencies. However, all too often they come up with the same stuff, and, as sure as eggs is eggs, there were a lot of things that I recognised. The people I met name-checked life sciences, they name-checked space and they name-checked AI. Those are all their growth sectors. So far, so good—you could replicate those anywhere else and, indeed, they were slightly surprised that Scotland was focusing on the same things, too; they had not realised it. What was interesting was that one of their domains for growth was the trust economy, and

that replaced two things: financial services and what we would probably normally put as tech. I thought that that was quite interesting, because it was identifying something that is more essential. Rather than how you are doing things, it is what the underlying point is.

You are right. I think that there are some elements there about Scotland and trust. What are the things that you think make Scotland a place that can focus on trust and how could we build a trust economy in Scotland? What might that look like if we wanted to outcompete Singapore on that?

Steve Aitken: First, as someone from Scotland who looks at Scottish people and is always surprised at how much of a mark we make in history for the population we have, I think that we should absolutely aim to have as massive an impact as we can. When it comes to the trust economy, the easiest thing to look at is banking. There is a massive amount of trust involved in putting all your assets and money in an institution; you have to trust that it is not going to take your money away. However, that is true of everything you do, because as soon as you exchange and as soon as you buy something, you are trusting who you are buying from. That is the essence of trade.

I am amazed that Singapore brought that up. If you get that as a brand for the country, you solve so many problems with trade with other countries. If you are generating that trust internally, you have less in-fighting and you have much more positive economic benefit for everyone.

How do we do that? I think that we do that by showing the positive benefits of generating trust, how to create trust and giving people stories around that with examples. A brilliant thing is that the most looked-at person from business on the planet is the perfect example. He was massively trusted and then did something that was not seen to be moral, so that trust vaporised. It vaporised completely and everyone's mind changed. We have the best example to point out at the moment to say to people that, if you are building a business and it is going really well, which I am sure he thought his was, there is always a place to fall to. If we can get that trust and keep it, we can perhaps do a lot more than we would have otherwise.

The Convener: Fantastic. We have definitely run out of time, although we have certainly not run out of questions. I thank both our panels for validating my degree choice, given that I am a philosophy graduate. With that, I thank you both for your time this morning. It has been extremely useful. We have a lot to think about and I am now worried about how we will pull this together in a single report. We might have to use some AI ourselves to do that. Thank you so much. I draw the public session to a close.

12:22

Meeting continued in private until 12:34.

This is the final edition of the *Official Report* of this meeting. It is part of the Scottish Parliament *Official Report* archive and has been sent for legal deposit.

Published in Edinburgh by the Scottish Parliamentary Corporate Body, the Scottish Parliament, Edinburgh, EH99 1SP

All documents are available on
the Scottish Parliament website at:

www.parliament.scot

Information on non-endorsed print suppliers
is available here:

www.parliament.scot/documents

For information on the Scottish Parliament contact
Public Information on:

Telephone: 0131 348 5000

Textphone: 0800 092 7100

Email: sp.info@parliament.scot



The Scottish Parliament
Pàrlamaid na h-Alba