

DRAFT

Economy and Fair Work Committee

Wednesday 12 November 2025



Wednesday 12 November 2025

CONTENTS

	Col.	ı
ARTIFICIAL INTELLIGENCE (ECONOMIC POTENTIAL)	1	
· · · · · · · · · · · · · · · · · · ·		

ECONOMY AND FAIR WORK COMMITTEE

31st 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)

THE FOLLOWING ALSO PARTICIPATED:

Claire Baker (Mid Scotland and Fife) (Lab) (Committee Substitute) Steven Grier
Professor Mark Schaffer (Heriot-Watt University)
Heather Thomson (The Data Lab)

CLERK TO THE COMMITTEE

Anne Peat

LOCATION

The James Clerk Maxwell Room (CR4)

^{*}attended

Scottish Parliament

Economy and Fair Work Committee

Wednesday 12 November 2025

[The Deputy Convener opened the meeting at 09:30]

Artificial Intelligence (Economic Potential)

The Deputy Convener (Michelle Thomson): Good morning, and welcome to the 31st meeting of the Economy and Fair Work Committee in 2025. My name is Michelle Thomson, and I am the deputy convener. Our usual convener, Daniel Johnson, is unable to join us today. We have apologies from Sarah Boyack, and a former committee convener, Claire Baker, is standing in for her—welcome, Claire. Lorna Slater has also given her apologies today.

This is our second evidence session on artificial intelligence. I am delighted to welcome our witnesses: Steven Grier, industry adviser and former director at Microsoft; Professor Mark Schaffer, professor of economics at Heriot-Watt University and fellow of the Royal Society of Edinburgh; and Heather Thomson, chief executive officer, The Data Lab.

I will begin with some opening questions. Following that—you are probably aware of the format—I will introduce colleagues to focus on areas of interest to them.

Last week, we had quite an optimistic session, in which we talked a lot about opportunities. Today, I want to start by asking about risk and getting a little more meat on the bones. Last week, I asked about the black-box nature of generative AI—we do not know what is going on in there. Today, I would like to get a sense of what our witnesses see as the critical risks in the area, particularly economic risks and risks to the public sector. I invite Professor Schaffer to start—for obvious reasons.

Professor Mark Schaffer (Heriot-Watt University): They may be obvious.

The list of risks is long. The risks are not just economy-wide; there are also risks for specific lines of work and specific businesses. For example, speaking as a university lecturer, I can say that artificial intelligence is not making our lives easier when it comes to education. We now cannot rely so much on the methods that we were using for assessing students, so we have to

change. We also have to teach our students how to use Al appropriately.

There are lots of challenges, and they are not going to be limited to education. You might have seen reports—last week or at the beginning of this week—that someone has written a bot that enables people to come up with responses to a planning exercise. That makes it easier for people to engage in participatory democracy, but it will also create issues with how we manage that input from the population.

Going back to the experience of the beginning of the dotcom boom, the internet, big tech and so on, regulators would say that, in retrospect, they did not act quickly enough, which meant that we ended up with an industry that is dominated by very large companies that were able to ensconce themselves pretty early on without serious constraints. This time around, regulators are more aware, but it is going to be hard to change what is happening.

I could keep going. Should I?

The Deputy Convener: Yes.

Professor Schaffer: Okay. The use of Al agents is going to pose problems and challenges for businesses, because businesses will want to know that they are dealing with a genuine person. For example, someone might use an Al bot to make reservations at three or four places and then they will go for the one that they like best. They might not put in that effort themselves, but they might be happy doing it if they are using a bot. That will pose issues for businesses when they get requests.

There are also risks for individuals who do that sort of thing, because they might be asked to supply credit card or bank account details and people will need to trust that they are using Al agents that are not going to make big mistakes.

The Deputy Convener: I can see what you are saying. There is clearly a great deal of potential and I suspect that, as we go through the conversation, there will be more examples of what might happen. You have already given us some useful examples that we might not have thought of.

Heather Thomson, can you answer the same question?

Heather Thomson (The Data Lab): The issue of trust that Professor Schaffer has just mentioned is one of the challenges with regard to public services as well as the private sector. The issue of transparency is important. The convener mentioned black boxes. It is hard to gain people's trust when they do not understand what is happening to their data.

A lot of the conversation about Al seems to focus on people's questions about what is happening to their data, but there is less discussion of the opportunities. Without pivoting the conversation away from risk, because that is an important issue, I would say that, if we focus on messaging around privacy, the general data protection regulation and so on, but do not also communicate to people that the use of this technology could change their lives for the better—for example, by getting them into a hospital bed for a life-saving operation more quickly—we will fail to use our messaging to develop public trust.

Not enough is said about real use cases and the benefits of the technology, but that discussion has to be very well balanced with the issue of risk, and the main issue in that regard is that of transparency, as people need to understand what is happening with their data.

One of our challenges with that is technological literacy. We are asking people of all ages to engage with systems that they do not understand. Young people are well placed to engage with these systems, but many elderly people who have a problem and need to speak to somebody do not know what to do when they are faced with an Al agent or a chatbot. There is a huge risk that those elderly people may just not bother asking for the help that they need because they do not know how to use the system. We need to ensure that services are working in the way that they were designed to work.

Of course, there is a risk in relation to young people. Although they are more tech savvy, there is a risk that they are becoming overreliant on Al technology and are not developing critical thinking skills because they use Al systems and just take whatever they get back. Rather than challenging the response or developing the ability to problem solve, they say, "We do not need to think about that; we can just ask Alexa." There is a risk that, if we introduce these systems from an early age, people might get to a point in their lives where they do not have the necessary life skills.

The Deputy Convener: Steven Grier, do you want to come in now?

Steven Grier: I would have liked to have done last week's optimistic session.

The Deputy Convener: I think that we will get on to that aspect.

Steven Grier: First, I echo Heather Thomson's point about the challenge regarding critical thinking, particularly for youngsters—when I think of the concept of my 15-year-old son saying "I will ask my AI", I do not imagine that that develops critical thinking.

Given my background, you might expect me to flip the question slightly and say that the biggest risk comes from the fact that, as a country, we do not move fast enough or adopt new approaches fast enough. Instead, we pontificate, we prevaricate and we do not move, and so we get left behind.

From an economic perspective, there are risks to be managed. I am sure that some of them were covered in last week's session-indeed, they are very broadly covered in about four or five different news articles this morning. One of them is the impact of AI on certain types of job, especially the jobs that young people predominantly go into-I am talking less about the service industry and more about the entry-level jobs in industry, which involve rudimentary rules and an administrative focus and are the jobs that are likely to be impacted by Al. That is a risk for the country, especially when we consider that fair work is one of the risks that we should address as soon as possible. I do not think that, this morning, I am at a stage at which I can suggest what we can do about it. However, I imagine that, as the conversation evolves, that is what we will get to.

Building on colleagues' responses, other risks involve poor governance, bad actors and an overreliance on Al in inappropriate circumstances. There are many ways in which the adoption of this technology can go wrong, as has been the case with every technological advance in the past 120 years. It is important that we have the right guardrails, the right operating environments and the right people using the technology.

It is also important that we trust the suppliers of the technology. The likelihood is that with an economy that is based on businesses that are small and medium-sized enterprises—SMEs make up 90-plus per cent of business activity in Scotland—most Scottish organisations consume their Al rather than build it. Do we trust those who provide that service? As businesses, when we contract with that service provider, are we happy that the legalities and the administration around that AI are what we would like to see? Scotland has huge potential in that regard, because we are very well governed and are a safe environment. Our marketing of Scotland as an Al nation should be built around that solidity and trust. We are well positioned to deal with the risks that are posed, some of which have been very well articulated by my colleagues.

The Deputy Convener: I can see that you are correct with regard to what you say about society and trust. We can take a view on the issue, but we are an advanced economy in all ways. Of course, there will be bad-faith actors, but we have seen that in relation to other areas, too.

My opening question was about the uncertainty around black-box generative AI, where it will be much harder to track what is being done and capture some of the risks-some of which could be insidious, depending on how we are populating the systems. We know that there are concerns about biases being built in-I understand that that is one of your areas of expertise, Professor Schaffer—but do we have a good enough sense of the known unknowns, in that respect? I do not think that anybody can say that we totally understand the situation, because of the exponential rate of change, but are there enough people who are worrying about the possibilities and doing the thinking about them?

Professor Schaffer, you are inclining your head, so I will bring you back in.

Professor Schaffer: The last part of the question is hard. Are there enough people working on the possibilities? I do not know. It is fair to say that a lot of people are.

Speaking personally, I am not massively worried about the fact that these systems are black boxes. They are incredibly complex things consisting of something like 5 billion parameters, and they inherently cannot be analysed in the way in which we would analyse computer programs in the distant past-that is, the previous millennium, when I learned my computer programming. Instead, they have to be analysed in a way that is more akin to how we deal with animal health and safety. Cattle are really complex biological beings, but we manage to deal with them. I do not think that the fact that we have to look at Al models as objects into which we have to feed inputs and that we characterise based on what comes out is inherently new or terribly challenging in principle. That is my personal take.

Heather Thomson: To an extent, your question goes back to the understanding of the system. From our perspective, as the Al hype continues, the foundational message is about the importance of the data that is used and of understanding the outputs that we will get. A lot of the bias and the harms will be caused as a result of the data going into these systems.

09:45

At The Data Lab, we have spent the past year conducting a piece of research on closing the skills gap in Scotland. We have considered what the skills gaps are currently and what they will be in the future. We surveyed more than 500 businesses of all sizes across the public and private sectors, and some of the results that have come back are very interesting. With regard to the challenges of what goes into those systems, fewer than a quarter of business leaders in the private

sector say that their organisation has a data strategy. That varies significantly by business size. Some 46 per cent of public sector leaders say that their organisation has a data strategy. When we go further into that research and start to look at equality, diversity and inclusion initiatives, we see that the organisations with data strategies are the ones that understand the importance of diversity in building these systems. However, only a small percentage of organisations have such a strategy; many people are just experimenting, which is where a lot of the problems will come from.

The Deputy Convener: Thank you. Stephen Kerr will ask questions about our next theme.

Stephen Kerr (Central Scotland) (Con): Steven Grier, you very quickly got to saying, "We need to go faster"—that is what you think we need to do. Just break that down for us: what does "go faster" mean for Scotland?

Steven Grier: Al is probably used in three ways currently. The first is conversational Al; the second is that a business will almost customise its Al for a business use or purpose; and the third is essentially the massive one, which is looking at biochemistry or predictive analytics in healthcare, for example. You can look at it in those three ways. First, we should provide the support and guidance that small business needs to comfortably adopt Al, and we should focus on productivity, because that is the most quickly and easily attainable benefit from Al.

Stephen Kerr: But we have been working on productivity in that sector of the economy for a very long time and there have been all kinds of initiatives, yet none of them has really made a dent, so what are we going to do differently in order to bring AI to that set of businesses?

Steven Grier: At the risk of solutionising it, the first thing that I would do is ensure that those organisations have access to people with the skills—I do not mean skills in creating AI; I mean skills in using Al. If they do not have that talent pipeline coming in, we need the ability to train them, and to skill or reskill them, through programmatic initiatives that are easily available and preferably free. We want people in a small business to go in and say, "I think that there is something that we do that can be improved by Al"—by ChatGPT, Anthropic, Copilot or whichever form of Al they choose. If they do not do that, those businesses will continue to act as they did before and they will potentially become uncompetitive. Smaller businesses will penalised more than bigger ones, because bigger businesses have the resources and the money to develop AI more quickly, and then we create that inequality again. Some might say that business is business, but one of the things that we need to do about that productivity question is to act fast and

make sure that people can quickly access the skills they need so that they can progress. Once they are in that world of improving their productivity through AI, they will, we would hope, progress to bigger and better things using the technology.

Stephen Kerr: But is there not a problem that is even more basic than that? Many businesses will see these tools as a way of facilitating existing processes and ways of doing things, because perhaps—and for understandable reasons—the business owners and operators and business leaders might not have the vision as to how they can re-engineer their entire business. Simply making courses available to them, introducing them to tools—

Steven Grier: That is a great point. One of the biggest challenges that we have in Scotland is the leadership confidence to make decisions like that. If I may, I will pivot slightly to the public sector, although I will come back to SMEs in a second.

Stephen Kerr: I want to get on to the public sector, so that is fine.

Steven Grier: One of the biggest challenges that we have is that people in leadership do not feel comfortable with the pace of the technology. If I were focusing development and skilling assets—whether that be universities, colleges, apprenticeships, or training courses—leadership is where I would go first. If we have leaders who do not feel comfortable in making the bold step forward with this technology, we will not move and we will lose competitiveness.

If I were sitting here with piles of tenners, that is where my money would be going. Public sector Al has been my passion for the past 10 or 15 years, so I feel quite strongly about it. We should give leadership the confidence to make bold decisions and to move beyond pilots—someone is bound to have said that we have more pilots than British Airways. We need to bring those pilots to life in both the public sector and business. How do we get them from pilot to production? Quite often, pilots fail because we are not bold enough to roll them out across the whole country or across a whole business. We have pockets of utter brilliance in Scotland—absolutely world-leading brilliance.

Stephen Kerr: In the public sector?

Steven Grier: In the public sector, yes. Take the work that Dr Gerald Lip is doing on breast cancer screening; in my opinion, it is world leading.

Heather Thomson: Another example is the Al colonoscopy programme, which has been running now for three years, I think. That started out in NHS Highland as a small project looking at how

we could transform colonoscopy, moving away from the traditional method to a smart pill that is swallowed—it is much quicker and much less invasive and leads to much earlier detection. That started to be rolled out through a ScotCat programme but then stopped. There is evidence of the impact of that but, again, it is in a pocket.

Stephen Kerr: Why are there pockets? Why are there those bright spots of excellence?

Heather Thomson: A lot comes down to where the funding comes from. Another example is our data skills for work programmes. We were very fortunate through the Edinburgh city deal to receive money to run a data skills for work programme. Part of that was a data and Al upskilling credit scheme. It was built on the Singapore model, where every citizen was given money for upskilling. We had the money, we built our data and AI skills framework—this is where the papers come from—and, over the past four years, we have upskilled 1,800 people at various levels to support them in retraining and moving. The programme focus for the was underrepresented groups and those whose roles were at risk of redundancy as a result of Al. The programme has been hugely successful.

Stephen Kerr: How are you measuring it?

Heather Thomson: We can track where those people go and what employment they end up in. The programme allows them to access roles. We talk about needing to upskill and retrain. How do we approach job displacement? By having upskilling and retraining opportunities for those who will be displaced.

Stephen Kerr: What about the point that Steven Grier made about leadership? He said that he would go first to leadership. You are rightly focusing on something that you have done that has borne fruit in terms of upskilling but, if there is a leadership blockage here, and if there is a cultural reservation about making positive decisions about the adoption of technology, reprocessing and re-engineering the way that we work, we will not get the productivity gains that we are looking for.

Heather Thomson: There might be two separate points there. I absolutely agree with what Steven Grier said. That is why, at The Data Lab, we focus on leadership in our executive education, because, from all the research and market testing that we have done, that is where we identify that there is a gap—there is also a gap in governance, with boards. There is a huge need to educate leaders and give them confidence. It is the same in schools. We talk about schools, and we think a lot about the pupils, but we need to focus on the teachers. I would say that that is one point.

As our discussion moved on to why are things not happening-pilots and small pockets of funding-I gave the example of the data skills for work programme. That was regional funding for Edinburgh and the south-east region. As a result of the success of that model, it was rolled out into Tay cities, through their city deal. Although it is great to have the opportunity to launch these initiatives as pilots through regional deals—there are regional variations—and they help us to get a better understanding of what is going on, they stay in the regions. There is no reason why such programmes should not and could not be rolled out nationally. We have proved that they are a success, so how do we move them to a national roll-out? We and others have shelves full of examples of where that has happened: money is invested, we run the programme for six months, and then there is no follow-on funding, so it stops.

Stephen Kerr: So short-term funding is an issue.

Heather Thomson: Yes. Some of the policy recommendations that came out in the recent paper were also about looking at multiyear funding for national programmes.

Stephen Kerr: I am still trying to get my head around the issue of changing the culture. We got on to the public sector very quickly; I am particularly aware of the challenges that we have with public sector leadership culture.

Steven Grier: I would just add a point on the financial side of progressing boldly with Al. The expectation should always be that the business case stacks up. If it does not, why on earth would we be doing it? We need to prove the gains. A challenge that I would give back to the public sector in particular is that business cases that are built on productivity and not on cashable savings do not progress anything like as fast as they should. For example, in presenting a business case to the public sector, I might say that a particular Al feature will free up two weeks of a person's time. I pulled that figure from recent United Kingdom Government testing and piloting of generative Al-it was Microsoft's Copilot platform. The area in question will save two weeks a year, according to the pilot study. In my experience, when we have said to people in the public sector, "This is going to save you a load of time" they do not give that anything like the weighting that is given to saving money—that is, "I spent this last year; I do not spend it this year." Rightly or wrongly, that dominates business case processes in the public sector still. If we go back and say that we are freeing up a week of a doctor's time across every doctor in Scotland, we should be jumping up and down to get that rolled out, without question. For whatever reason, we do not. Again, I am not going to solutionise it here,

but I will say that, for the public sector to make the best use of this technology, it is going to have to look at productivity and cash-flow savings with parity of esteem.

Stephen Kerr: I think that "solutionise" must be a lovely Microsoft word. I have not come across it.

Steven Grier: It has taken me a while to cast aside my—

Stephen Kerr: Microsoft is still deep in your heart. I want to tempt you to be a bit more solutionising in the way you approach this, because I certainly think that we as a parliamentary committee would be interested in specific actions that we could take in the public sector to jump-start a change in that culture. So please feel free—solutionise away.

Steven Grier: I will be brief, because I am sure that my colleagues have much better-informed views. I am at risk of being overly simplistic, and I am sure that a chief financial officer of a public sector organisation would batter me for saying this. What if you expand on Gerald Lip's work in the national health service, where he has productivity and efficiency gains? He was essentially measuring in public sector consultant hours. Maybe this exists, but at no point have I seen someone take the average salary of a consultant, extrapolate it into the productivity and give me a cash sum at the end-so, if we save 10,000 hours of consultant time with this technology implementation, the number that we are going to use in the business case is 10,000 times the average hourly cost of that resource in the public sector. I have never seen that. It may be erroneous, and it may be completely disregarded in terms of a public sector structured business case, but I have never seen an attempt to do that and I feel that we should do that.

10:00

Stephen Kerr: Through your business experience, where have you seen that being used outside Scotland?

Steven Grier: It is not even a matter of looking outside Scotland—outside the public sector here, it is used in commercial business cases all the time.

Stephen Kerr: That is in commercial terms, though—I am asking about where it is used in the public sector. The commercial world has a completely different mindset about resource management and return on investment. In my experience, that does not prevail in the public sector too much, but it might in pockets.

Where have you seen that approach being adopted wholesale in the public sector? There will

be lessons to be learned from it that we could bring to Scotland.

Steven Grier: I do not know, but it would be well worth finding out.

Stephen Kerr: Okay. Heather, with your experience, do you know of a similar situation?

Heather Thomson: I will just clarify something with Steven Grier. In your example—

Steven Grier: I am saying that when anyone goes to the public sector with a business case that suggests that, for example, we will save 1,000 hours a year, I never see those hours. It would be easier for the public sector to consume a business case based on such a saving if we could put a numerical value on the 1,000 hours.

Controversially, one challenge is that we never want to talk about people displacement, which is obviously a factor. Is it actually a saving if those hours are not calculated? My preference is to look at the hours, especially if we are talking about a situation in the NHS. Where is the redeployment? Where are the reductions in waiting lists or times? If we needed to put a financial saving on those, or a financial redeployment within the NHS, I would start to look at how we calculate those cases.

Heather Thomson: I give the example of a project in the NHS that we have been involved in along with the William Quarrier Scottish epilepsy centre, which measured brain signals and analysed data to predict when seizures might happen. By implementing AI technologies, the team was able to reduce the data analysis time from 12 hours to 3 seconds. That is huge, especially when we think about all the other work that could be done in that way, which could improve the number of patients seen, early detection, the time it takes to have diagnostics from X-rays, and so on. At the moment, someone who goes for an X-ray will get their result six weeks later, because of the time that it takes for that process to happen.

The NHS absolutely is the place to look for examples of where such work is happening. Time saving makes a massive difference there and is clearly accepted as being an impact.

Steven Grier: One metric that we often see being used in the NHS is the cost of hospital admission. I have seen that being included in Al use cases.

Quite an old example that also happened in Glasgow, and which pioneered early intervention for patients with chronic obstructive pulmonary disease, was an early Al data-led project. In its construct was the concept of saving individuals from having to visit accident and emergency departments. The team replaced that A and E cost with a much cheaper alternative, which involved

proactively sending a COPD carer to a patient when they thought that they would be in trouble. Instead of waiting for patients to come to A and E, they would look at a red-amber-green report that said, "That patient is in trouble and is likely to come to A and E." Their response would be to send a COPD carer out to the patient's house, or to adjust their oxygen remotely, so that they saved them that visit.

Obviously, the focus in that project was on the patient experience, which became far better. However, there is always a cost attached to a visit to A and E. I cannot confidently tell you that that cost was included in the particular case of that project, but, again, it would be well worth finding out. However, the concept very much covered how costly a visit to A and E was versus the much more cost-effective option of proactively sending someone out to the patient, which was also better for them. Therefore, such an approach does exist in places.

Stephen Kerr: You are talking about calculating cost savings on the basis of budgets, but massive opportunity cost savings could be made by, for example, optimising the productivity of the health service.

Steven Grier: Absolutely.

Stephen Kerr: That should also be part of the business case, so you are really advocating that we overhaul our public sector procurement approach to take on board the business case.

Steven Grier: Absolutely. Our approach needs to move with the times. There are other challenges in the finance world—around how we pay for and consume services, for example.

Stephen Kerr: I am being smiled at by the deputy convener, so I will stop there. [*Laughter*.] However, it is clear that there are lots of other issues that emerge from what you are saying.

The Deputy Convener: There are, and we might have time to come back to them. Thank you very much.

I will bring in Kevin Stewart for a quick supplementary question before I bring in Gordon MacDonald.

Kevin Stewart (Aberdeen Central) (SNP): I am really interested in the conversations that have taken place about the work of Dr Gerald Lip.

Earlier this year, I lodged a motion in the Parliament to highlight Dr Lip's contribution to large-scale clinical trials using AI, particularly the GEMINI project—Grampian's evaluation of Mia in an innovative national breast screening initiative. I believe that his work has led to 12 per cent more cancers being detected than has been the case in routine practice, which is quite incredible.

If we look at the cost aspect of all that, we might view it not only in budgetary terms but in human cost terms. Surely such an advance is great for all of us. The human cost of an early diagnosis is better for the patient, so the human cost of their illness is likely to be lessened. Looked at from the perspective of health economics or societal economics, getting a diagnosis and treatment more quickly, which is likely to lead to better outcomes, should also mean that that person can be fit and healthy again and get back to being productive.

We should be talking more about the Gerald Lips of this world. Why are more universities, hospitals and health boards not looking to create the type of appointments that Dr Lip has, in leading the use of artificial intelligence in clinical practice? Why are we not moving such activities on more quickly? Why are we not using Gerald Lip as something of an evangelist? Why are we not hearing more about that kind of work?

Steven Grier: Every point that you made there is absolutely right. I do not have heroes, but, if were to have one right now, Gerald Lip would be right up there.

In a clinical setting, it is right that extreme caution is used. Gerald Lip is a strong ambassador for that. I stand to be corrected, but I think that he is now leading a similar exercise that is certainly Scotland-wide if not UK-wide. I believe that, quite rightly, he has been given what could be called an evangelist's role in using AI in cancerscreening technology across the NHS.

You asked why we are not doing more of that work. Again, I go back to the need for confident, bold leadership, but we also need to be assured of what we are doing. It is an incredibly delicate area, so we would want to be really sure. I am certain that Gerald would love to see things moving much faster than they are.

The numbers that Kevin Stewart stated came from a subset of scans that looked only at breast cancer. If we were to extrapolate that approach to what I would unaffectionately call the big four cancers, and look at it in the context of all the national screening programmes, we would see that the numbers could be absolutely phenomenal.

To go back to the question, I wish that we could be bolder. However, obtaining funding is always a challenge, as is achieving effective leadership.

Kevin Stewart: I am sorry to interrupt. Obtaining funding is always a challenge—there is no doubt about that. However, the reality is that, alongside the lessening of the human cost, the savings here could be huge. If we were detecting illnesses and treating people more quickly, the outcomes would be likely to be much more

positive, which would mean that a person could become productive again more quickly.

Steven Grier: I have one final point. How do we, as a Government and a collection of public sector entities, manage to pass the savings from Gerald's study back into the economy, to build an even more powerful business case? Right now, such work is happening in the NHS. There is no assignment of any kind that says that that person will be more productive sooner, they will be back in the workforce, they will not be claiming disability benefits, or whatever. In my view, there is no ability for us, as a nation, to look at nationwide cost justifications.

Kevin Stewart: That is because the UK is very backward in such regards, whereas health economics in other parts of the world is much more sophisticated in that it looks at the whole-life aspect of treatment. I could give other examples, but I will not because they do not relate to Al. However, such examples exist in other areas. Dr Lip is the lead on artificial intelligence in clinical practice at the University of Aberdeen. We should have somebody doing a study of his work from the health economics side, which would build an even bigger case for advancing the use of Al in our health services.

I notice that Heather Thomson wants to come in.

The Deputy Convener: I could also see Professor Schaffer shaking his head vigorously at one of your—

Kevin Stewart: I missed that—I am sorry.

Professor Schaffer: It was not that vigorous. I do not think that health economics is particularly backwards in the UK at all. The concepts required for measuring those things—

Kevin Stewart: Let me explain myself, Professor Schaffer. I am not slagging off health economists—I know a number of them, and I might get myself into trouble. I am simply saying that politicians do not look holistically at all the work that health economists do.

Professor Schaffer: That is fair. We could use international comparisons to judge what the UK gets in return for the money that it spends on health.

You will be able to tell from my accent that I am not originally from here. If we look at QALYs—quality-adjusted life years—and at what the United States spends on health versus what it gets back, the return is absolutely appalling. The US is at the technological frontier, so a lot of money is going in and such research is being rolled out at a faster pace. However, at a holistic level, it represents really poor value for money.

Here, the return is pretty good based on the resources going into research, but, based on international comparisons, those resources are pretty low. Therefore, although you are getting really good value for money, the money that goes in is just not that high compared with what happens in other countries.

Kevin Stewart: Thank you for allowing me the supplementary, convener.

Gordon MacDonald (Edinburgh Pentlands) (SNP): Good morning. Forgive me, but for a wee while there I thought that I had entered the health committee's meeting. [Laughter.]

As I was listening to what all of you said, I noted that there are questions about trust, transparency, bad actors and overreliance on AI, and that this area needs confident leadership. I also noted a comment made at the beginning of the meeting about how, at the time of the birth of the internet, regulators did not move quickly enough. What should the role of the Government be now, as far as AI is concerned?

10:15

Heather Thomson: Much of the current conversation in this realm starts with infrastructure, which is a complex issue and not one in which I would claim to be an expert. We could talk about the AI bubble, skills and adoption, but we can only go so far without having the right infrastructure in place. There seems to be a lot of conversation, but not a lot of action on infrastructure progression.

We can look at what is happening in other countries—and even elsewhere in the UK, as a result of the UK opportunities action plan and the Al growth zones. We are holding our breath while we wait for an announcement on an Al growth zone for Scotland, but that is not the be-all and end-all. A growth zone would be one area, but other data centres are looking to put down roots in Scotland. There needs to be an accelerated conversation and plans need to be put in place around that.

However, I say again that I am not an expert. There might be things happening behind the scenes that we are not aware of. There is the skills aspect, which we have already spoken about. If we are talking about what we need to do, I would say that we need to go faster and we need to take the issue seriously.

We talked about international comparisons and about how funding will always be an issue. I saw that at last week's session the committee heard about levels of funding for AI adoption programmes. When we compare those with what is happening in other countries and regions, it is

clear that we need to be taken more seriously in this area. Therefore, we need to act more seriously.

Steven Grier: Clearly, you have a very tactical and long-term strategic role in the public sector. You hold sway in that particular environment. If we, as a country, would achieve the saving of lives and the associated increased productivity that we need, gaining encouragement from the Government would be massive. I say that the Government should get out of the way when it needs to; it should be supportive when it needs to be; and it should provide funding for business cases that are based on factors such as productivity and cash savings.

Next, we need to think about what we want to be as a country. Do we want to be an adapter, by which I mean one that simply encourages our businesses to adapt to AI and use it to remain competitive? Alternatively, will we be bolder and say that we will be a global AI centre? If we choose the latter, what will our specialist area be? The Government will have to look at it and ask. "If we have four, five or six key economic areas"—in the same way that this committee will have, and to use one such area as an example—"given that we consider ourselves to be global renewable energy experts, we are a superpower in that area. Would we look at ourselves right now and say that, from the perspectives of academia and business, we are the AI experts on renewable energy?" I suspect that the answer would be no. The next question should be: why not?

If we want to focus on particular economic areas, fund them proactively and say that our country is incredibly good at them, we will have to expand them, including in the infrastructure of renewables. Behind that, we will have to build an education, academic and skills environment that shows not only that Scotland will produce renewables as a global leader but that it will build an AI renewables global centre of excellence and ask what that would look like.

We will not cover every single area of AI, but there are areas in which Scotland is immensely strong. The expectation is that if we could build that, very soon we would have international interest in what we do and we would have the potential to attract foreign investment into Scotland.

Gordon MacDonald: We have talked about the need to support AI and to help to grow its use, and about making us a centre of excellence. However, is there a need for regulation, given that there are issues with quality, trust and probably even ethics? Professor Schaffer, since you are the person who put that thought into my head at the beginning of the meeting, I will ask for your view on that.

Professor Schaffer: Yes, I did. About six months ago, I gave evidence to the Constitution, Europe, External Affairs and Culture Committee and the idea of regulatory frameworks for AI came up. Speaking practically, globally there will be a small number of large, wide-ranging regulatory regimes. For example, there will be one in the US and one in the European Union. The UK might be able to forge its own regulatory regime. Scotland on its own is too small for that, so it will not happen here.

One big issue is that things are moving really fast, so preparing a regulatory regime needs to take that into account. I have to mention a passage from a book, because it is just so wonderful. "A Brief History of Artificial Intelligence: What It Is, Where We Are, and Where We Are Going", which was published in January 2021, listed six problems that it described as being "nowhere near solved": understanding a story and answering questions about it; human-level automated translation; interpreting what is going on in a photograph; writing interesting stories; interpreting a work of art; and-human level general intelligence. We have now achieved five out of those six—we are there already.

Kevin Bryan, an economist at the University of Toronto, pointed out that the process of designing the new EU AI regulation was begun before that book was published. Therefore, the EU was designing a regulatory regime way in advance of massive change. What we want in a regulatory regime is one that is not extremely detailed and prescriptive but gives regulators a lot of flexibility to deal with new problems as they arise, because that will be needed.

Many issues have arisen already. For example, OpenAI is being sued by a family in the US because ChatGPT gave—well, actually, I do not think that I need to go into the details. Here is my point. We need regulation, and it is important that we have it, but it has to allow for new technology that needs room to experiment, but, at the same time, be flexible enough to deal with new issues as they arise and which we cannot foresee.

Gordon MacDonald: How do we get things right so that we do not stifle innovation? Has any country already gone down that path and started to get the regulation right?

Professor Schaffer: I do not think that it is a question of which country; the issue is supranational. There is the US, but it has a federal system so individual states—for example, California—are trying to introduce their own AI regulations. The UK is probably big enough to implement something on its own and, as far as I understand it, the direction of travel is one in which there is flexibility. It is not as prescriptive as the system in the EU where I think—I ask Heather

Thomson to correct me if I am wrong—they specified something about the limits to computing power.

Heather Thomson: In the regulation?

Professor Schaffer: Yes, in the regulation. That is kind of crazy.

Steven Grier: One challenge would be the structure of the Scottish economy. As I said at the start of the meeting, if an economy consists predominantly of SMEs the likelihood is that it will consume AI as a service rather than build it. Such a situation relies on putting our trust in the governing structures that have been agreed to the nearest macro level—whether they be global, as Mark Schaffer said, or at EU or UK level—and the providers of that service will adapt to the regulatory regimes in which they operate. That is both challenging and useful if everything is going well; it is perhaps more challenging if it is not.

It would be difficult for Scotland to effect that, given the scale and size of the businesses that own those services. An element of trust is being placed in broader governance regimes. In my view, we are well governed. There is a lot of awareness about the safety of AI, and about trust and transparency, and we currently see those issues being worked out right across the governmental landscape.

Heather Thomson: That is where we have an opportunity, too. We talk about Scotland wanting to become a leader in AI. What it actually wants to do is become a leader in responsible AI, building upon the work that has been done to date by organisations such as the Scottish AI Alliance. In the governance that we see in place at the moment we have a platform on which to build. Some of the conversations in which our team are already involved include people in the US who are looking to Scotland as a route, because they want to be part of an ecosystem that works responsibly.

Gordon MacDonald: Okay, I will leave it at that. Thank you.

Claire Baker (Mid Scotland and Fife) (Lab): Thank you for being here this morning.

I suppose that the average person's understanding of their engagement with Al involves trying to work out whether the video on their phone is Al or not. Yesterday, on the radio there was a story about reducing animal testing and Al being part of the solution, but, this morning, there was another report on the radio about the music industry, the pressure that it is under from Al and the concerns that exist in that area.

The example of health has already been cited. It feels as though, in health, AI is a tool. I think that most people would understand it as a tool in the sense that we would usually recognise a tool—it is

able to analyse and provide information more quickly and to be reliable in doing so. However, for most people, that is only part of the story of Al. Could you say a bit more about where you think that there are legitimate concerns, whether for sectors or for individuals? Is improving public understanding part of the solution, or is the public's level of distrust or fear legitimate?

Professor Schaffer, would you like to respond? You mentioned the issue of students and how difficult it is to manage their use of Al in universities.

Professor Schaffer: We should distinguish between different types of Al. The point that Heather Thomson made at the beginning about data is important. Al depends heavily on the data that it is trained on. In science, we control what goes in, and what we get out of it is clearly defined. In areas such as animal testing, molecular structures or designing drugs for the health sector, the position is straightforward and, in my view, the dangers are not huge in the same way that they are when we talk about uncontrolled use of public data.

Of course, the general public will not think about such distinctions, but they are right to be worried about AI use. This morning, I saw that someone on X—formerly Twitter—asked whether the 2020 US presidential election had been stolen, and Grok, which is X's in-house AI, and is under the control of an individual, said, "Yes, it was." That is really dangerous, and everybody should be concerned about that. However, that is different from animal testing or designing drugs.

Claire Baker: Heather, in the work that you do, do you distinguish between the ways in which, broadly speaking—there are probably more than two of them—Al is used? When we think about the economy, are we just more focused on it as a tool? Next week, we will hear from a panel that will include musicians, who might have more to say about copyright issues and so on. Is your organisation more focused on how Al can be used by businesses as a tool, or do you also engage with broader issues around interpretation or the potential to mislead?

10:30

Heather Thomson: Absolutely. An issue that we have spoken about today, which was also mentioned at last week's session, is the need to educate people. That is part of our focus. We want to help people to understand not the value of Al but the value of their data, and what is the art of the possible when it comes to their data. We will not even go near an Al system until we know the answer to questions such as, "What data do you have?", "What are you using it for?", "Is it in

order?" and "What is the quality of that data?" Fundamentally, one of the biggest challenges is that people do not know that information.

We have recently been in situations in which we have had conversations with people who have told us, "We don't have any data in our organisation." When we said, "Really? Don't you collect this and this?", they said, "Oh yes—we do." We told them, "That's data." That is the level of literacy that we can be dealing with on the part of people who are using AI tools, which is terrifying.

There are concerns about, for example, bias if people do not understand how such systems operate. It does not matter whether it is a tool or something such as ChatGPT. Last week, I went into a classroom to talk to schoolchildren about the issue. Imagine a class survey that involved asking everybody, "What's your favourite lunch?", but only the data from the girls was put through; the boys were not asked. That is the sort of data that can go into such systems. There can be a lack of diversity. There can be gender bias. There can be all sorts of things that bias results in. Unless the data set that has been used to get those results is challenged or queried, you will never know that.

For me, one of the key points that came out of last week's session was that we should never use only Al tools to take a final decision. The best and most impactful way to use Al, and the safest way of adopting it, is to have a human in the loop. They can augment the role of Al quite significantly. When it comes to image diagnosis in healthcare, there will be further checks in place. A human will not have to go through every single image, but it will not simply be a case of, "Computer says yes" and that is it. The danger, which there is legitimate concern about at the moment, is that, in the absence of the necessary education, people are just accepting what comes out of Al systems and are not challenging it.

Claire Baker: Professor Schaffer mentioned the election and the use of misleading videos. Business does not happen in a bubble. People who run businesses and SMEs do not isolate themselves in their businesses. Like everyone else, they use their phones, so they will see external influences. We have talked about badfaith actors. Businesses need to think about how Al could be used negatively against them, perhaps through comments about their business. That happens more often these days. How can we increase people's knowledge and understanding of that?

Steven Grier: In part, that is an extension of—

Claire Baker: The Government is coming up with an AI strategy, although I think that it has been delayed until the spring. There is a subgroup of industry members. The process is very

business focused. Do you think that that is the right approach, or should the Government look at the broader impact of AI on society?

Steven Grier: I would say that it should definitely do the latter. In responding to that question, I will mention a couple of areas.

At the risk of making this argument seem not as important as it is, Grok is simply like *The Times* was in 1925. In those days, if people read something on the front page, they believed it. They chose their newspaper. We are nearly in that world. That is an extreme example, but we have always had to take care when it comes to the information that we receive. I go back to the point that was made in the first five minutes: the absence of critical thinking is a danger. At this point, we are in long-term philosophical discussions about Al. The ability to see, believe and act is dangerous. It was dangerous 100 years ago, and it is dangerous now.

Claire Baker: But, with the pace of change, it has become more difficult to tell whether something is true. Previously, we would see an animation and we would be able to tell that it was not real, but now that we can see an image of an actual person doing something, it is hard to believe that what our eyes are telling us is not true.

Steven Grier: We also now have the concept of Al policing Al. The only way that we can police that is by having alternative Al engines look at something and verify whether it is real. A great deal of technological progress has been made in relation to, for example, fingerprinting of images. Such technology has been around for a while, but it has relevance in this age. Through a reverse image search, it is possible to look up an image to find out whether it is real or involves an actor from a different country.

A related issue is the fact that, if we do not build trust, we will not get the best out of the technology. There is a trope that I have wheeled out for 20 years that says that a person is 50 times more likely to fill in a form to get 10 per cent off in a big retail store than they are to give their information to the NHS to save their life. I used to give the example of my wife in that context, but it became unacceptable for me to say that my wife was more likely to do it.

We face a cultural challenge in that regard. I cannot speak for other countries, but that is certainly the case in the UK. As Heather Thomson said, we need to have a balance through equality of input. We cannot make national decisions on predictive healthcare if we do not have everyone's data—I am sorry to take you back to the health committee again, Gordon. If a certain demographic of people do not trust the system

and are excluded from solutions as a result, we will have failed to get the best out of a tool that has the potential to be incredibly impactful.

Heather Thomson: To pick up on Claire Baker's question about the AI strategy or the AI action plan and the focus on business. The Data Lab was involved in the AI strategy that was developed five years ago. The vision that came out of that AI strategy was for Scotland to become a leader in ethical, trustworthy and inclusive AI. A lot has happened in five years, but, at that point, there was a sense of, "What is this thing? Is it coming to take over everything?" There was a need to educate the public to settle them, without even thinking about the huge opportunity for business that existed.

You asked whether the Government's approach should be wider. It absolutely should be. Since the initial strategy was launched, everything has evolved. Last week, a remark was made about the fact that the existing AI strategy does not cover energy and renewables. I would argue that, at that time, we were not thinking about the impacts that AI would have or about the need for data centres and infrastructure. It is absolutely the case that the new strategy needs to involve more than just business, but we are quickly realising the opportunity that exists, and it is important that we move now to grasp that opportunity or we will miss it

In my view, the original strategy was very strong on the ethical governance side of things, and it should not be diluted in that respect. I have always been of the opinion that ethics should never be a bolt-on—it should not be bolted on at the side but should be embedded and ingrained in everything that we do.

I go back to the point that I made earlier to Gordon MacDonald. We have an opportunity to be known worldwide as a nation that believes in responsible and reputable Al. With the new strategy, there is an opportunity to build that in and to focus on economic growth, but not to the detriment of civic society or of inclusion. We need to think about how we grasp that opportunity while minimising the risk. That will involve us bringing in a whole number of things that were not in the initial strategy.

Claire Baker: Can I ask one other brief question that is linked? We are an ethical nation but a small country. Can we have an ethical approach to AI, given that a lot of the content is not generated in Scotland? How difficult is it for any country—and we are quite small—to say that it will be an ethical AI provider, producer or user when so much of the content comes from outside?

Heather Thomson: I will give a brief answer and let others come in. I think that it is about

having frameworks in place and about people understanding. I feel a bit like a broken record, but it is a lot about understanding the systems that you are using and understanding where your data is going, which country it is going to and what the frameworks are there. It is also about producers. You have the different parts of the model with AI, and we talk about homegrown AI. It is about understanding what frameworks need to be put in place to ensure that there are standards that are followed at that point.

Professor Schaffer: Al inputs are turning into a commodity that is supplied, and there are markets for them. It is possible even for a small country, on a smaller scale, to regulate the use of commodities. Influencing the supply side is a different kettle of fish, however, because there you have scale issues, international regulation and so forth, and it would be really challenging for a small country to influence that. However, on the use of AI, its being commoditised, regulating that and putting in frameworks and ethics, I agree completely with Heather Thomson that that is not a bolt-on. That is feasible.

Heather Thomson: There is ambition. Going back to the infrastructure piece, if Scotland gets itself to a point at which we have the data set, we have the power that we need and we have the data centres so that, when the organisations are growing, the data is staying in the UK, it then becomes much easier for us to control and manage. Ultimately, that is the long-term vision. How do we, as a nation, ensure that we are not exporting our power down to the UK and that other people's economies are not benefiting from the work that has been done in Scotland?

Murdo Fraser (Mid Scotland and Fife) (Con): Good morning. I would like to ask more about the issue of workforce and skills. However, before I do that, I want to go back to an earlier discussion about the public sector and productivity. A story on the front page of The Herald newspaper today that caught my eye is very relevant. It says that the head of the Scottish Public Pensions Agency is being summoned to this Parliament's Finance and Public Administration Committee to answer questions about the delays in providing compensation remedies for hundreds thousands of public sector workers. That is to do with the compensation for unfair discrimination.

The SPPA has been given 18 months to give pension remedy statements to those affected, but it has missed two deadlines, meaning that retired people are being locked out of their entitlement, and it is costing taxpayers millions more in interest, which is currently charged at 8 per cent a year. Why on earth are such processes still being done manually by the SPPA, in this day and age, when we could be using AI to do them?

Steven Grier: You would not be using Al there, Murdo. You would not be using digital, either, I suspect, in a lot of those cases. You would be looking at paper-based returns, claims being made through an antiquated process and an approval loop that is horrible. I am guessing here, because I do not know the facts of that case, but it is a massive frustration in the public sector, which is, in my view, the most promising area for huge productivity gains, for sure, because we have situations like that. We have huge paper-based processes that we are reliant on but that should be digitised in the first instance, and then the data should be assembled in the right way. That is a data challenge; it is not an Al challenge. We clearly do not have the data in the right place, or we do not trust the data that we have assembled to fix that.

You would imagine that, in—I am trying to think of the right word—the panacea that is AI driven, that process would take a matter of days rather than a matter of months if the data was in the right place and was correct. You made that case fairly eloquently there, and *The Herald* has made that case indirectly today, I would say.

10:45

Heather Thomson: They may be digital systems, but there are many legacy systems in the public sector and, in order to implement AI, the data needs to be in a machine-readable format. That, in itself, is a challenge.

Steven Grier: I have some sympathy, because the records will go back years and years, and the investment in digitising them would not have been justified. That is a hump that we may have to get over. Many medical records in the UK are still held in paper format, although Scotland has an advantage there. In fact, Al has been used to digitise health records south of the border, to enable them to come into this world.

Professor Schaffer: Maintaining systems is incredibly expensive. That is where government is different from business. Government has an obligation to maintain these things for much longer periods than private business does. You only have to look at the Windrush scandal, in which lots of records were just tossed. Digitising them would have been incredibly expensive and, at the time, maybe hard to justify, but look what happened.

Murdo Fraser: Thank you. That was bit of an aside, really, but it is quite interesting, and, as it was on the front page, I thought I would ask you about it.

I want to talk about AI changing the nature of work and skills. I know that we touched on this earlier, but I want to probe a little bit further. I am

looking at a report that Microsoft did on 17 October. Steven Grier, I do not know whether you have seen this.

Steven Grier: It was after my time.

Murdo Fraser: Okay. Well, it is called "Working with AI: measuring the applicability of generative Al to occupations". It is quite a detailed and interesting report that looks at occupations in which Al has the highest applicability and, therefore, is potentially the largest threat to people working in those sectors. Among the top five, we passenger attendants and representatives of services, as you might expect. I will not name them, but I know of at least one local high school that I have been to that prides itself on its vocational training, which provides youngsters with the skills to work in call centres, doing salestype jobs. That will be reflected right across Scotland. Do we have the right skills training available, either in schools or further up the chain, to equip young people—or people at any point in their career-with skills that will not be made redundant by AI? If not, what should we be doing instead?

Professor Schaffer: I will take that question, as I work in the education sector. Al skills will be needed. It is not just that people will be replaced by Al, because, as colleagues were saying, you need that human element. You need people to be trained in how to use Al responsibly and sensibly, which will bring productivity gains, too.

I think that the question should be: are we equipping our students at school level and at university level with the appropriate skills to manage these new tools? It is starting to happen, but it is taking a while. In the education sector generally, the first perception of it was that it is a threat to the way we teach, as I mentioned at the beginning, and in some sense it is. However, as educators, we have a responsibility—and not just a responsibility—to send our students out into the world knowing how to use these tools. It is a moving target, so it is really hard, but it is happening—maybe not fast enough, but it is happening.

Heather Thomson: I need to be careful in how I answer the question.

Professor Schaffer: Did I get myself into trouble?

Heather Thomson: Prior to being CEO at The Data Lab, I led its skills programme for the previous seven years, so I have seen a lot. I also have two young children: one in primary school and one in secondary school. One of the things that I cannot understand is the talk about having pockets of opportunity and innovation when there is inconsistency in what is being taught to students. It is not even regional—schools in the

same town are inconsistent in the opportunities that they offer, based on the way that the curriculum is set through curriculum for excellence, and the freedom that they have to teach in order to get the outcomes that we need. I will not even try to remember the number of secondary schools in Scotland that do not have computing teachers, but it is a massive problem. It is also about the curriculum and how we change it and bring industry in to help the schools to understand what skills they need to teach. Technical skills are important, but more important are critical thinking, problem solving and team working, which can be taught in any subject.

From a skills perspective, I was talking about ethics earlier and I have always tried to be an ambassador for how we embed these skills into everything that we do. Prior to working with The Data Lab, I was employed in another part of the university and worked on how we help to increase the quantitative skills of social scientists. Quite often, social scientists are the people who go for those subjects at university because they may not be good at maths. They are not good at maths so they go towards the more social subjects. They spend a lot of time focused on qualitative research, and then, when they come to work in policy, they start to understand how you evidence that. Where are the numbers? Where are the statistics?

From a schools perspective, not having the computing teachers that we need is a huge problem, but how do we equip teachers in every subject to understand that data is everywhere, numbers are everywhere and AI is everywhere? How do we embed that into the learning so that it just becomes part of everyday life? It is not even about working for a job now; it is about your daily life. How do we teach citizens to be able to function in today's life?

Quite often in schools, things are taught without a context. Last week, I went into a school to talk to a group of nine and 10-year-olds about data and AI. They had spent a term doing data, and I asked, "What have you learned about data?" They said, "Bar graphs, line graphs, spreadsheets." I asked, "Do you understand why? Do you understand the importance of these skills in everyday life? Do you understand the opportunities that await you, some of which we don't even know about yet?" I did not talk to them about data; I talked to them about the opportunities, and the context actually brought those subjects to life.

How do we excite people about this? Given the opportunities that are available to young people now, through graduate apprenticeships and modern apprenticeships, and in the fact that you can go to university and get paid to do a job at the same time, I am not sure that I would try to

convince any of my children to do a standard university degree. There just seems to not be that consistency, and it seems to be out of date.

One of the challenges is that things evolve so quickly. We have worked with universities around the curriculum for data science and data engineers, and we have brought in industry advisory boards to understand the needs. However, given the time that it takes for that to then influence and get into the curriculum, by the time it is there, it is already out of date. That is a huge challenge as well.

Steven Grier: Exactly as Heather Thomson said, that is where we rise up a level, take the four constituent pillars of curriculum for excellence and say that, if we get that part right, those skills can be used in the world of AI.

The computing challenge is horrible. There are some things that we really have to fix there. I would be interested in Willie Coffey's view. There are some challenges that we have around the image and perception of computing in schools. I am an advocate for the abolishment of the word "computing", because I do not think that there is a bigger turn-off, particularly for young females coming into the industry, than the image of computing. Skills Development Scotland and other bodies have done work to improve that and to change things. I can give you some interesting anecdotes about changing the name of a class from "computing" to "digital design" and having quadruple the number of applicants for that particular class—just through some semantics and some marketing around what is a hugely exciting and creative profession, which has now been augmented by the potential for Al.

Your original question asked what we should do about the fact that we are training people in schools to be in call-centre jobs that might not actually be there in future. As I said at the start, I do not profess to offer a solution to that. We could be reskilling, retraining and repivoting education in the use of AI so that, when someone does go into a job, the first thing that they are thinking is, "How can I make this better? How can I use AI to make what we do quicker, faster, more profitable and less impactful from a harm perspective?" Exactly as Heather Thomson says, we need a way to move more quickly. We should not be sitting there with out-of-date teaching materials and content.

It goes back to the issue of teaching the fundamentals of computing versus most people simply using a service. When I build an application now, I will go to something like Copilot—others are available—I will assemble it and I will ask AI to create my customer service tool for me. Then I will plug some agentic AI into it, so that that can deal with the requests when they come in, pivot on the basis of the data that it gets, do the next one and

pass it on to the next AI agent. I will be able to do that massively quickly.

We, as a society, have to decide whether we need to teach the fundamentals of computing. I think that we do, but we also have to be cognisant of the fact that the world of application and system creation has changed beyond recognition.

Murdo Fraser: What you are saying is fascinating. Did I see that you recently joined the board of the Scottish Funding Council?

Steven Grier: I did, although I must put my hand up and say that I am an unaffiliated resource for you today.

Murdo Fraser: That is fine. It is good to know that the SFC is—I hope—leaning on your knowledge and expertise in this area.

I was interested and a little concerned to see in the Microsoft report that the top five occupations on which AI will have the biggest impact include historian and author. I declare an interest in both categories. Is there a risk that human creativity will be squeezed out by AI? If the report says that historian and author are jobs that will be squeezed out, what does that mean?

Professor Schaffer: If you are asking whether there are really serious implications for creators in the creative space, the answer is yes—absolutely. We are even seeing music that has been generated by AI being popular. It is going to be harder for people in certain sectors to maintain a living. The creative sector is the one that we are talking about, but there are others. However, will they be entirely squeezed out? I do not think that we are at that point yet.

In response to your previous question, I add that everything that we have said applies to university education as well. We expect all our university graduates to be literate, and we should expect them all to be digitally literate, too, whatever their discipline.

11:00

Heather Thomson: This week, the University of Southampton published a series of essays about the impact and implications of AI in universities. As part of the research, one of the essays has been written by ChatGPT, and it is there for people to compare, review and comment on. That might interest you. I will share it after the meeting.

Murdo Fraser: Thank you.

Steven Grier: I do not know the context of the Microsoft report—I have had a self-imposed media blackout for a month in my camper van—but I would be interested to read it. I think that Al's potential as a tool must be so exciting for an historian. I cannot think of anything more exciting

than building patterns and correlations between data that we have not realised before. I am not meaning to put a positive spin on what I suspect is a concerning aspect of the report, but I would be focused on and excited about that.

I think that you will get better answers to the question. I do not mean that those were not brilliant answers, but I imagine that, if you were to have artists and creatives on a future panel, they would express their concerns very vocally.

Heather Thomson: From the augmentation perspective, Al can pay dividends in helping people to write. How do we structure and review writing and how do we proofread? Al can take someone through the writing process. However, when it comes to Al replacing writers, I am not so sure. I do not know that I would want to read something written by Al. To me, it would not feel like it contained emotion and I do not think that I would get as involved in it, but that is just my perception.

Steven Grier: I have a 15-year-old son and I see him using AI as a research tool to pull in information that he would not otherwise have. If we presume that, as Mark Schaffer said, we can trust in what we read and we are not being deceived, that is really exciting. I cannot decide whether to be alarmed or super happy about it, but he is accessing information that he would not ordinarily have, and he is building correlations, coming to assumptions and being enthused. We had textbooks and then the internet, and he is now having things created for him by AI. It is a glib statement, but that is incredibly interesting and fascinating.

Professor Schaffer: I see what Steven Grier said about historians in my discipline. Economists are normally equipped with numeracy and digital skills. Economic history is a growth area, with more and more stuff being digitised. For economists who work in economic history, it is a fantastic opportunity. Historians who are digitally literate are going to do really well. There may not be demand for so many historians in the future. That is an entirely different matter. However, in terms of the productivity of historians, there are massive opportunities.

Willie Coffey (Kilmarnock and Irvine Valley) (SNP): Good morning, everybody. This has been a fantastic conversation. I will start where Steven Grier ended a wee while ago, on the subject of computing. How to encourage and keep girls in science has perplexed everyone for many years, and we still do not know the answer. It seems that, when girls transition from primary to secondary school, they lose interest in science, and it is as though computing becomes akin to the oily rag. Mechanics, engineering, software engineering and computing seem to turn young girls off.

Last week, Sarah Ronald talked about her company, where younger women really excel at data analytics and like that side of computing science, whereas the younger guys like to be the coders and the programmers. I do not know how true that is in general, but that is her experience. Is there a magic wand for how to persuade more young women to stick with computing? I think that the idea of calling it "digital design" is fantastic.

Steven Grier: Considering what we can do about that has been a passion of mine for a number of years, but there are better, far more qualified people than me to talk about it. For example, you could get Toni Scullion from dressCode to come in and talk about how she sees it.

It is partly a question of semantics and marketing, but when we consider the creativity that can be driven by the digital world, it is difficult to understand why the imbalance exists. We are obviously doing something wrong. The image of a spotty youth, who is usually male, gaming until their eyes are blacker than soot and being fed pizza under the door—I am deliberately being humorous—seems to present a challenge. That is sometimes the image, and we need to work to improve that. Some work can certainly be done there, but I am not the best person to talk about it.

Some really interesting work was done by Skills Development Scottish Scotland in the Advisory Apprenticeship Board's aender commission, which looked at levelling up both industries where young males were not coming in, such as care, and industries such as engineering and digital, where the challenge that we are discussing still exists. Some interesting work was done there from an apprenticeship perspective, and it is well worth reading around things that companies can do to try to fix some of that.

Willie Coffey: This is just an observation, but I remember showing my daughter a smartphone years ago and saying, "Wouldn't you like to be able to programme and control these?" She said, "No".

Steven Grier: People just want to use them.

Willie Coffey: Yes. For many people, the value is the functionality of the thing. They are not really interested in how it is designed, how it comes together or how powerful it can be as a tool. They just want to be the end user of it. I do not know how representative that is.

Steven Grier: A niche point is that, when we look at Minecraft, that particular game or ecosystem has a quite good gender balance. That is an anomaly. It may be because it is hugely creative, but it encourages just about everything that we want. Why do we have that up to a certain

age point but we suddenly do not seem able to develop from there?

Heather Thomson: I am a computer science graduate. I was one of three females in a class of 80, and by the end of the first year I was the only female. I moved into the industry, working in the utilities sector and then the offshore sector, so I have come through a very male-dominated environment. Having spent so much of my time in those environments, when I came into the university and The Data Lab and got more involved in the diversity in tech, it was actually a surprise to me, because I had been so ingrained and embedded within it that I did not really see the anomaly. As time has gone on, I have learned and reflected on the impact that that had on me and the way that I changed to make myself fit into the situation. It has been a really interesting journey.

How do we get more women involved? Last week, I attended the Scotland women in technology awards, and I have not been in a room that was more inspiring. A 25-year-old young lady from JP Morgan who collected an award is already sitting on governance and advisory boards. It reminded me of the old saying that you cannot be what you cannot see.

How do we spotlight women in the industry? Mentorship is hugely impactful. I have been involved in Dell's STEM aspire scheme, which takes in female undergraduates and college students every year and assigns mentors to them. It has been interesting to understand the levels of imposter syndrome; the challenges that people feel, even in walking into a room in those scenarios; and the cultures that are associated with the male-dominated sector, with people feeling awkward and like they do not fit in.

We talked earlier about the role of jargon. Terminology is used in school and university teaching, but as soon as we put in the words "computing" or "engineering", people just switch off. How can we create a more inclusive learning environment for people whereby they do not feel that things are only for men or only for women? Actually, they are for everybody.

Willie Coffey: That is a wise message.

Professor Schaffer: Earlier, we discussed digital literacy in schools and universities. If something is a requirement and everybody goes through it, it levels up the playing field, so that might help.

Willie Coffey: What I really want to talk about is ethics—it always creeps into the conversation; we get there eventually. How can ethics be embedded at the heart of the AI revolution? Can it or should it be? Perhaps it already is. Earlier, Heather Thomson said—I scribbled it down—that ethics

should be embedded in all aspects of Al. How can that be done?

Mark Schaffer, you said in your opening remarks that corporations grabbed the whole agenda and ensconced themselves early on. They were not thinking about ethical standards. They were thinking about profit, control and influence and all the rest of it. Can we truly embed ethics into Al, or must we rely on, for example, governance or regulatory measures in order to throw some kind of protective blanket over it?

I would be pleased to hear your thoughts on how we do that. You started this, Heather, so you can go first. I add that I have never seen an ethical computer algorithm yet.

Professor Schaffer: Oh, really?

Willie Coffey: Have you? Professor Schaffer: Yes.

Willie Coffey: You can tell us about that later.

Heather Thomson: That goes back to the conversation about bad actors and human influence. There is input to AI systems, and ethical considerations stem from the people who influence them. Are frameworks in place for that? Morality is a factor.

On getting the right balance between regulation and non-regulation and between ethics and innovation, it is very hard to achieve that without having in place frameworks and guidance, so those elements are important.

We have been involved in conversations recently about schemes that are almost incentives for people to use AI, whether that be access to the compute that they need. I am involved quite a lot just now in conversations with data centres, and one of the questions that I ask them is, "What ethics controls do you have in place? How do you know what your compute is being used for and how does that impact your reputation?"

For me, that is not just one person's responsibility; it is a whole chain of responsibility. It is about people coming together in the ecosystem to agree what it is that they want for Scotland. You will never get rid of the bad actors, but we, as a conglomerate, if you like, can come together and agree what guidelines we will follow and what we want Scotland to be known for.

11:15

From an educational perspective, we should teach ethics and governance from school age. That needs to be built in; it should not just be an add-on. Someone should not design a system then ask how to make it ethical, or do so at the point at which they are interacting with it. We need

to think about the problem holistically and take account of those important considerations. It is a very difficult question to answer.

Willie Coffey: That is why I am asking it.

Professor Schaffer: Examples of ethics in coding and in designing AI systems are mortgage applications and credit scoring. We have protected characteristics, and those go into databases. Should we train AI to make predictions that a protected characteristic is associated with a lower probability of repaying a loan? If that were to happen, people with a protected characteristic would be less likely to get a loan. That would not be ethical.

Ethics get built in straight away. If you train a facial recognition Al on a population that is predominantly pale, it will struggle when trying to recognise faces that are not. What do you do about that? You can train it on a representative sample. However, they are minorities, which means that some faces will still not be recognised.

Ethical considerations are everywhere when designing algorithms. That is why I was nodding vigorously when Heather Thomson said that teaching ethics is not a bolt-on.

Willie Coffey: Should it be more of a voluntary thing, with some sections of society deciding that they will engage in that way, or does there need to be an overarching framework that everyone should observe? You mentioned Grok earlier. Is there an ethical component to Grok?

Professor Schaffer: There is a wonderful internet meme of a little girl shrugging and saying, "Why not both?" It has to be all over. This is a Scottish Parliament committee and we are talking about the regulation of businesses. However, really, we are talking about everything. It is hard for a small country to influence the design of the algorithms that are used by the behemoths. It is a lot easier to influence the design of algorithms that are used by local businesses. It is also a lot easier to talk about educating people on how to use such things ethically and how to interpret them.

Steven Grier: We have talked a lot about readiness and skills. If we are going to be an ethical AI nation, we should embed that foundation as early as we can, as Heather Thomson said. I go back to my earlier example. If we are to be the AI centre of excellence for what we consider to be our key economic priorities, our reputation as a country should be that we are ethical in how we go about doing so. We control what we control. If we are going to create, for example, an NHS data set, we have control over that and our use of it. However, we might use a tool such as ChatGPT, Anthropic or Copilot to pull out that data, and we have to trust that we can do that properly.

We will control what we can and then ask for strong governance of the tools that the majority of us will ultimately use. Practically, that is as much as we, as a small country, can do.

Willie Coffey: Finally, a number of colleagues around the table mentioned education. What would be the direction of travel were an Al component included within education? The briefing paper from our friends at the Alan Turing Institute told us that 49 per cent of the time that is spent on activities in education could be better supported using generative Al tools.

I invite you to gaze forward and say what education might look like in the immediate future if AI should become more embedded as a tool for learners and teachers and for the activities that are traditionally used to engage people in education. What might AI become if it is deployed sensibly and ethically in the education setting?

Professor Schaffer: There is a lot of scope in education for that. I am speaking at the university level, but I guess that this works at the school level, too. Students learn well when they interact with teachers, but the level at which they interact can sometimes be basic. Nevertheless, they learn from the interaction. It is possible to automate that. Education tools are commercially available in which you can train an Al on a particular body of material and it assembles answers to students' questions. Those are right/wrong answers. However, there is also potential for interaction between the student and the Al. That is just one example.

There are also many challenges. As I mentioned, we are having to radically change the way in which we assess. We may have to go back to the medieval viva era, where people must discuss interactively with another human being what they have done in their work and their explanation for it. If that involves asking them, for example, how they assembled their body of work and getting them to explain what they used to do it, their prompts to the AI and how they stresstested it, I am totally okay with that. I am not sure about many of my colleagues, though.

Heather Thomson: That answer focused a lot on the student perspective. There is a need to support teachers, and Al can be a massive asset for them. We all know that teachers and university and college lecturers have significant workloads and are under a lot of pressure. How can we support them through Al to become more innovative in their teaching? They can use Al for writing lesson plans and for introducing ideas and contextualising some of that learning by using it to identify case studies, new activities and interactive activities, and for bringing lessons to life.

I talked earlier about the data skills for work programme, which is one of four workstreams as part of the data-driven innovation skills gateway. Another workstream examines data education in schools. Moray house school of education and sport at the University of Edinburgh leads on that. It has done phenomenal work with schools, particularly one in Midlothian, on how we introduce Al into education. That includes learning about some of the sensors and the micro bits, and helps pupils to understand how, for example, their school could contribute towards net zero by putting solar panels on the roof and what the cost savings of doing so would be. Excitement can be generated by, for example, getting data from NASA.

Willie Coffey: Steven Grier can have the last word.

Steven Grier: I do not have a lot to add. If we had AI as a research tool, we would be so excited, and the kids and the young people would be so empowered, by the challenges that AI creates.

If we go back to the core of what Mark Schaffer was saying, that is about the good old concept of pupils showing their working. We must ensure that young people in schools still understand that, and that they have not simply copied or generated their work. Even if they have learned from the process of researching on AI, can they explain the theory? How do we then assess that in a modern, progressive way?

Professor Schaffer: That is how it will work in the workplace.

Willie Coffey: Will the pupils of the future have their own personalised AI bots that look after their individual educational development journeys?

Steven Grier: In a way. I hope that AI will be used to spot areas of inequality within education and to spot attainment challenges so that we can then use that evidential data to say, "Here are the interventions that we must make". I am excited about the idea of linking that to a data set, linking it to demographic self-records and then being able to create a profile for that child that says, "Here is your ideal learning environment and your learning style". I am hoping that what AI will bring to education is an individual focus on the learner, which is attuned to them.

In my view, there is nothing wrong with how we have done it so far, but this gets discussed in education forums all the time: are we creating loads of little clones of various levels of ability by teaching them the same thing? The idea of personalised learning potential, driven by this technology, is hugely exciting for education.

Willie Coffey: Okay. That is absolutely fascinating. Thank you.

The Deputy Convener: Thank you all very much for your contributions. Before I bring the session to a close, I have one last, rapid-fire question for you all.

We mentioned earlier that the Scottish Government intends to produce its AI action plan and new AI strategy in the early part of 2026. What are the top three things that you would like to see from the plan and/or the strategy? We have had a very wide-ranging session here today, but what are your top three priorities? I know that it is a difficult question, so do not look away, anybody. Professor Schaffer, what are your top three?

Professor Schaffer: I will try not to look away. I am sorry, but it is very hard to prioritise something that is so—to use the jargon—high dimensional. There are just so many dimensions to it.

Maybe what I would be looking for, rather than priorities, would be the characteristics of the process. Things are moving so fast that I would like to see a strategy that is inherently flexible and which is regularly reviewed and revised. Maybe there could be an annual cycle for revisions. Coming up with a strategy and then sitting back for three years is doomed to failure, I think. That is two things—is two enough?

The Deputy Convener: Yes, thank you—l appreciate that it is a difficult question.

Heather Thomson: It is about understanding that the biggest risk is to do nothing; we need to act and we need to act now. As we have spoken about already, we need to move away from temporary grants towards investment in multiyear national skills planning, taking stock of the huge, nationally funded assets that we already have and looking at how we can support them to grow in scale, as opposed to starting again and having little pockets of activity everywhere. Bringing things together is a huge thing. Rather than having pockets of activity, we need to take a more coordinated approach so that this feels like team Scotland—it feels like Scotland's approach to AI and it does not feel as though it is diluted or broken up. We need to encourage stronger partnerships across the public, private and academic sectors to make sure that, when we are looking at innovation and skills, we are all on the same page and everything is relevant.

Steven Grier: You asked me to be brief, but I do not know how I will make this brief. I would love to see us create an economic environment that attracts AI compute capacity and capability into Scotland. We have not talked about that a lot today. How are we creating the economic environment that encourages a hyperscaler to come here?

I would be proactive in building confident Al leadership in the public sector. We need to build

on the amazing cases that we have so far and decide how we want leaders to think about the opportunity that Al presents in the public sector.

I would embed AI thinking in the world of employment, economics, apprenticeships and education. I would go back and look at every single apprenticeship right now and see how each one is potentially impacted by AI and how AI could be used better, so that we are not starting with a young workforce that is already behind those in other countries.

The Deputy Convener: Thank you very much to you all. That brings our public session to a close. I really appreciate all the time that you have given up this morning and all the information that you have given us.

11:30

Meeting continued in private until 11:44.

This is a draft Official Report and is subject to correction between publication and archiving, which will take place no later than 35 working days after the date of the meeting. The most up-to-date version is available here:

www.parliament.scot/officialreport

Members and other meeting participants who wish to suggest corrections to their contributions should contact the Official Report.

Official Report Room T2.20 Scottish Parliament Edinburgh EH99 1SP Email: official.report@parliament.scot Telephone: 0131 348 5447

Fax: 0131 348 5423

The deadline for corrections to this edition is:

Tuesday 16 December 2025

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



