HOUSE OF COMMONS

Science and Technology Committee

Oral evidence: Governance of artificial intelligence (AI), HC 945

Wednesday 8 March 2023

Ordered by the House of Commons to be published on 8 March 2023.

Watch the meeting

Members present: Greg Clark (Chair); Aaron Bell; Dawn Butler; Katherine Fletcher; Rebecca Long Bailey; Stephen Metcalfe; Carol Monaghan; Graham Stringer; Christian Wakeford.

Questions 174-262

Witnesses

I: Professor Andrew Hopkins, Chief Executive, Exscientia.

II: Professor Delmiro Fernandez-Reyes, Professor of Biomedical Computing, University College London, and Adjunct Professor in Paediatrics, University of Ibadan; and Professor Mihaela van der Schaar, John Humphrey Plummer Professor of Machine Learning, Artificial Intelligence and Medicine, and Director, Cambridge Centre for AI in Medicine, University of Cambridge.

Examination of witness

Witness: Professor Hopkins.

Q174 **Chair:** The Science and Technology Committee is in session. We are pleased to continue our inquiry into the governance of artificial intelligence. Starting off our session this morning, which is looking at the applications of AI in medical settings and medical science, we are pleased to have Professor Andrew Hopkins, the founder and chief executive of Exscientia, which pursues an automated and algorithmic approach to drug design and discovery. Prior to founding the firm in 2012, Professor Hopkins spent 14 years at Pfizer and in academia researching the use of data and machine learning in the pharmaceutical industry. Thank you for coming to give evidence today. Will you describe the work that your firm does, building on your prior career with Pfizer and in academic settings?

Professor Hopkins: Absolutely. When I used to be a member of the pharmaceutical industry back in the day, I was always very frustrated by the sheer cost and time that it takes to convert new science into new medicines for patients. This is an incredibly expensive industry. Much of that expense, and much of the reason we pay so much for drugs, is about the cost of failure, because many ideas do not work. So we looked at this, and I concluded it is very much an information industry. If we think about it that way, we can then think about how we can use the latest techniques in information technology, such as artificial intelligence and machine learning, to improve the way we handle the data and think about decision making. In fact, that is what the whole R&D process is in the pharmaceutical industry—it is very much around data and decision making.

Exscientia was set up as a truly AI-first company to consider how we can change the way the R&D process is conducted and how AI could be applied truly end to end, through the whole process of how you design the drug, how you develop it and so on.

What we fundamentally see, then, is that in creating a more overarching approach to how AI can change the approach to R&D, you also potentially end up with much better drugs. It is really about how you can design the right drug and select the right patient for that drug. Overall, we believe that can change how we operate as an industry.

Q175 **Chair:** Can you give us an example of each of those areas? How can the application of AI improve the process?

Professor Hopkins: I will give you two examples. The first is how AI can help us select the right patient and how we can improve patient outcomes.

We published a clinical trial in the journal *Cancer Discovery* last year. We believe this is the first time that an AI machine learning system has been shown to improve clinical outcomes in oncology. What we had developed was a system where we would take samples from patients with blood

cancer, do single-cell analysis—that means looking at individual cells—and then see the response to over 100 different types of drugs, which is far more than you could ever give a patient to trial. All the patients involved in the trial were basically late-stage patients. They had gone through five, six or seven lines of therapy, with very little left to treat them with.

We used the AI system to really analyse all the different ways all the different cell types—your cancer cells, your immune cells, your healthy cells—would respond, because you want to distinguish and to find a drug that will kill off a cancer cell, but not kill your immune cells. What we found, by using the complex biology of this, was that we needed AI systems to really analyse the single-cell details of all the different drug responses. We were then able to rank all these drugs and say, "For this individual patient, this particular drug is the one you should recommend to prescribe." That is truly personalised medicine.

The results of our trial were that, in the control group of about 15 people, within about a year, everybody was progressing. Within the group that was treated with drugs recommended by the AI system, around 25% of patients were still progression-free survival after four years.

This was looking at a whole range of drugs using the existing pharmacopoeia, as it is called—the existing medicine cabinet—and saying, "Which particular drug could be useful for them?" We would find, for example, that a drug that had potentially failed a clinical trial for a particular blood cancer indication was actually the right drug for that patient, because, in the clinical trial, they weren't necessarily selecting the right patients. That's often why clinical trials fail—it is often due to statistics, because you haven't necessarily got the right patients being selected for a trial.

So that had what is called an objective response rate of about 55% and a hazard ratio of 0.53. These would be excellent statistics for a new medicine, but this was a new process to help us get the most out of our existing medicine cabinet to help patients live longer, with progression-free survival for much longer.

Q176 **Chair:** Let's pause on that for a second. In our first oral evidence session, one of our witnesses said that this is not new—this is maths, basically. Isn't what you've described—looking at the effect of a particular drug on a particular patient—what medical research has always done? You analyse its impact. What is the difference that AI is offering in this?

Professor Hopkins: The difference is the complexity of biology that we are now analysing. Now, we are looking at individual single cells, not just gross amounts of tumour. We are looking at the statistics of hundreds and thousands of individual cells and the distributions of different responses. We are now also looking at how we integrate all the different omics data—single cell sequencing is possible, single cell transcriptomics, and the single cell phenomics of how we respond—and distinguishing between, "Is that a cancer cell? Is that an immune cell? Is that a healthy cell?" and looking at how we respond differently.

That's where computer vision and machine learning become incredibly important—in helping us distinguish this out. This is something of a scale that a human pathologist would never be able to explore at this level.

Q177 **Chair:** That's the point, isn't it? It is using data that is there, but you are saying there are limits to the cognitive or computational ability of human individuals, so the computing power allows more to be seen in the data. Is that a fair description?

Professor Hopkins: Absolutely, and the relevance of that as well. Many times people would do tests in laboratories on cell lines and mouse models, which actually have very little relevance to human disease. What you find when you can do these kinds of detailed analysis, which actually only machine learning is capable of allowing us to do, is that you can truly do it with a personalised approach. We have been talking about personalised medicine for nearly 20 years. What we are starting to see now is the first trials of it truly happening—that is, this particular drug for that particular individual at this particular point in time.

It is an enabler, and it enables us to take the intellectual framework that we have had for 20 years but, now, put it into real practice. That is what machine learning is allowing us to do. The complexity of the data that we are now gathering and that we are capable of gathering—as you know from your other investigations into the genome revolution and how cheap it is now, potentially, to gather genomic data—and the sheer wealth, depth and breadth of that data, mean that it is beyond, as you said, human cognition to understand how we would analyse it. The machine learning approaches we are using actually allow us to do that.

Q178 **Chair:** One final question from me, and then Stephen has a supplementary question. You mentioned that the control group was 15 people. That seems quite a small group. Why so small?

Professor Hopkins: The whole study was over 70 people altogether. It was our first investigational study, actually. We have just set up—we just announced yesterday—an expansion of this trial approach with Charité hospital in Berlin, taking it further. This was the very first publication of this approach. What you would expect to see now is more trials and further development of this technology as we take it forward. But this technology is not just used for the patient of today. This is the second example: we also use it—

Chair: Before you go on to the second example, Professor Hopkins, let me bring in Stephen, who has a question on what you have said so far.

Q179 **Stephen Metcalfe:** It goes back to one of the earlier points you made, which was about the use of failed drugs—drugs that failed previously at trial, perhaps because they were too specific to a certain cell type. What do you do with that information once you have it? And does that mean that there is an argument to revisit lots of other drugs that have failed at first trial or second trial stage, because they could potentially be personalised? The other part of that is this: is there a huge cost involved with this? We always assume that personalised medicine is very

expensive medicine. Is that the case?

Professor Hopkins: No. The higher response rate we can obtain fundamentally creates a much more beneficial economic situation for healthcare systems and for patients. As I said, the real cost is the cost of failure—the cost of waste. If you can match the right drug to the right patient, you actually start to reduce that waste, for the benefit of patients and for the benefit of society.

Yes, there probably are a number of drugs that have potentially failed because they were in the wrong patient groups and because of the way they were tested. The more we can precisely define which patients should be prescribed a particular drug, the more, overall, we should start to see in time, potentially, success rates start to improve. That a key feature of how you can use the technology I have just described—not just, potentially, to help a late-stage cancer patient today, but to help us design better clinical trials and select better patients for a trial.

So this is what we are doing with the new drugs that have been generated by AI and that are now heading into the clinic—novel medicines. We now take the wealth of omics data that has been generated and really help distinguish between what we believe are very sophisticated biomarkers, as they are called—gene signatures—to say, "We believe that that patient group is much more likely to respond than this other patient group."

This goes beyond the first generation of genomic medicine. With the first generation of genomic medicine, as you might be aware, you might have a particular mutation in a cancer and you can design a particular drug for that mutation. But in many cases, even if you are targeting that particular patient group, you still might only get a 30% response rate; there is obviously something else going on in biology beyond just that single mutant. This is what the more sophisticated approach is for. But it is only really enabled by the use of machine learning to interrogate and understand the wealth of data that we now are able to generate from patients and so on.

Q180 **Stephen Metcalfe:** Just a quick follow-up or a clarification perhaps. Where drugs have previously failed and have therefore not been taken forward, it is presumably in the interests of the companies that developed those to revisit their bank of development research. Are they starting to do that? That is what I'm trying to get to. Has what you have said motivated people to revisit drugs?

Professor Hopkins: Potentially—

Q181 **Stephen Metcalfe:** Only potentially.

Professor Hopkins: Potentially; this is absolutely at the cutting edge. These results were published only a few months ago. It is certainly possible that any company would want to see that. Could the money they have invested into a generation of that IP potentially benefit patients? I believe there is a new opportunity to be created around the right match-

making, so to say—identifying the right drug for the right patient. I believe that will fundamentally help us improve the odds in the clinic and so on.

Chair: Brilliant. Carol, did you want to come in?

Q182 **Carol Monaghan:** Again, this is just a slight clarification. If a drug has previously failed a trial, how do you persuade a regulator to allow that to be used again in your trials?

Professor Hopkins: If it has failed for efficacy, and one has generated new data showing that it is actually an effective medicine, then that passes through the normal regulatory process as usual. If it has failed for other reasons, such as safety, then it is a very different situation to consider. Here we are really talking about efficacy. Does a drug work for patients? We are not yet discussing other issues.

Q183 **Carol Monaghan:** So these are not drugs that failed for safety reasons, just for efficacy.

Professor Hopkins: Yes, exactly. Did the statistics for a trial fail because we could not find the right people?

Carol Monaghan: Thank you. That is fine.

Q184 **Chair:** That is a very important clarification. Professor Hopkins, you were going to describe your second example.

Professor Hopkins: Yes. The second use of AI is around how it can make the processes of creating new medicines much more efficient. You might be aware—it is in the public imagination right now—of the concept of generative AI such as ChatGPT, VALL-E, and Midjourney in the artistic department. In terms of the use of generative AI in chemistry and drug discovery, Exscientia has been developing for nearly a decade now, since we first published our research in *Nature.*

What we see now is that, if we work with our expert drug hunters, using these generative algorithms basically to help us design molecules towards very specific objectives—to make the drugs potentially safer—it is about precision engineering. You should think of a drug, not as a little white tablet, but as precision-engineered technology. The atoms arranged in that molecule are the result of potentially billions of other decisions that could have been made about what it could be. In terms of reducing that research space down, AI allows us to search it and reduce it far more efficiently.

To give an example, we have now developed about seven prototype drugs—several of them in patient clinical trials, several others heading towards clinical trials in what are called an IND-enabling studies. What we see now is that, on average, we have made around 90% fewer compounds—about 250 compounds per project, compared with around the 2,500 compounds that would be usual in a pharmaceutical industry drug optimisation project. That has implications for cost and for time and, potentially, we believe, in terms of much more precise molecules being designed, as well. We are now seeing repeatable evidence that the use of

computational approaches—AI approaches—in drug design can potentially help us create a more efficient drug discovery process.

That is what is really exciting about this. What you do then is think about combining that kind of generative approach to a potentially more efficient approach where you have drug discovery, with the approach I described a few minutes ago about selecting the right patient. That is what we are trying to do: bring all those together into an end-to-end process. That is what I see as the multiple benefits of the application of AI across that whole industry spectrum.

Q185 **Chair:** Good, excellent. It is being more precise about how you can target drugs to people, and drug discovery to design drugs based on problems that you want to address.

Professor Hopkins: Absolutely. It is an information problem of data and decision making.

Chair: I see. Good. Let me turn to my colleagues, starting with Rebecca Long Bailey and then Aaron Bell.

Q186 **Rebecca Long Bailey:** You mentioned efficiency in the development of new treatments, and also that there was an emerging evidence base. Obviously, this is a new area of science. Why do you think that the larger pharmaceutical companies have been reassured enough to make a significant bet on AI in this field?

Professor Hopkins: When I founded Exscientia and we first did this work, it was well before AI was fashionable. We had a deep belief that this was the way forward for industry. The evidence base that has been accumulating—drugs moved on to the clinic, the kind of clinical trial that I described earlier called EXALT-1—is now giving the pharmaceutical industry increased confidence to think about making these investments.

They are thinking about investments not just in research, but across the whole pipeline, particularly in commercial and marketing as well as in R&D. There is a concerted effort of investment. If you speak to companies such as AstraZeneca and GSK in the UK, you will see that they have been hiring quite a number of data and AI scientists.

One advantage that a company like Exscientia has is that for us it is not plan B; it is plan A. In fact, it is our only plan. We believe that this is the way to reinvent the process. The true advantages of AI are about not just adding it to an existing process, but thinking about how we can reinvent a process. For example, if you think about finance and the use of computers in finance, we did not actually put robots into the high street to dish out cash as a bank teller. We completely reinvented a process where all your banking is now done on your mobile phone via an app. It is a way to think about how you might re-imagine and re-engineer entire processes. That is the exciting thing, and that will be some of the interesting competitive interplays between the smaller insurgent companies like Exscientia and the larger companies.

It is important to remember that there is a lot of collaboration in this ecosystem. We collaborate with a large number of big pharma companies like Sanofi and BMS, and that is important. It is about each company bringing their different skillsets. Big companies are very good at commercialisation and late development. Companies like Exscientia are very good at discovery in the early development stage.

Q187 **Rebecca Long Bailey:** The issue of data protection is a huge concern in the emergence of AI in healthcare. How can AI systems be provided with the data they need to improve in a way that innovates while also maintaining adequate data protection?

Professor Hopkins: That question is at the heart of what this Committee is investigating. Particularly in healthcare, we are already in a very well-regulated industry, which is as it should be. We not only have the regulations of healthcare, the ethics of running clinical trials, and the safety regulations for developing pharmaceuticals, but we have all the data governance and privacy issues that we need to take into account. We have myriad overlapping regulatory frameworks to work within. What we are trying to do here is to think about how we can make a very expensive industry much more efficient in the long run. That, I believe, is for the benefit of patients, in bringing more and better medicines forward, and also for the benefit of healthcare systems and Governments ultimately.

One of the things to consider is how we in the UK we can make advances so that AI is seen as an enabler for the benefit of society. The risk is not applying AI in these places, rather than applying it. You already have a very strong framework in terms of data governance, confidentiality and ethics frameworks within those multiple regulatory frameworks that I described. At the same time, the UK has also been an amazing place for generating data. Look at what we have done with our sequencing initiatives and with UK Biobank. Those have been world leading in providing large data environments that can attract innovation to the UK.

Q188 **Rebecca Long Bailey:** One of the concerns is that even if you try to regulate data, as we do here in the UK, and you remove names and details from the records you are using as part of a dataset, certain indicators might still point to a particular individual—the date on which they had a heart attack, for example—and that would flag them up and breach their data. How would you try to overcome those sorts of data protection breaches?

Professor Hopkins: That is beyond my area of expertise, so I will talk about how I think about the systems. It is entirely possible to think about the use and ethics of the data that we generate, particularly as you think about NHS hospital data going forward. There have been a number of initiatives trying to do this in an ethical fashion, and they are to be supported. To get the framework right, think about it in terms of not just a piecemeal approach with small individual companies maybe taking a few hospital trusts and using their data; consider it as a national strategy. It is actually one of the huge potential data resources that the UK has. It is

almost equivalent to oil in the Saudi desert; you still need to mine it or drill it out to get the value.

The unified healthcare system we have in the UK and the data records that go with it are potentially of huge value to the UK. Therefore, if you can think of the right regulatory framework that would allow you to be comfortable with that and think about it as a national strategy—seeing it as a key data resource—then this is a very important potential source of wealth for the UK. It could act as a honeypot event for innovation and to attract companies to the UK to use the data.

Q189 **Rebecca Long Bailey:** One other issue that is certainly of concern is bias in datasets. How would your company try to overcome bias?

Professor Hopkins: Of the datasets we are looking at, which look at individual patients, the bias we think of is ensuring that we truly understand that, for example, the particular sample of cancer we have is representative of someone's entire cancer. We think of bias in a very different and individualised way. In terms of societal bias, as you alluded to, the key is to think about how we create larger datasets, and then ensure that we are sampling across the entire population correctly. It is about ensuring that underrepresented groups are included in healthcare records and how we exploit that. That is a key thing.

There are potentially ways of designing those systems to incorporate and understand bias. One of the key things we deal with is the fact that drug discovery is both a big data problem and a small data problem. We know very little sometimes about starting things, and therefore techniques—what we call active learning—have been vital. We have to ask ourselves, what bit of data should we collect next to actually improve our models? Active learning techniques could be one way to think about how we understand the bias in our models, and then deliberately to go out to seek particular datasets to improve the model in a certain way.

Q190 **Rebecca Long Bailey:** My final question is on transparency. Can the AI systems used by businesses like yours ever be fully transparent and explainable, or are they by definition black boxes?

Professor Hopkins: No, they are not by definition black boxes. It depends entirely on the architecture one is building the systems on. We like to understand the provenance of a prediction. When chemists are looking at a design and asking, "Why is it suggesting that we design a molecule in a particular way?" we can trace that back and reply, "Because datasets say that for this particular target, or to avoid this particular side effect, you have this design". That explainability, as you could call it—connecting the dots between a prediction and the data that led to a prediction—is vital for understanding in science, as much as it is for the general public.

It is not inherent that any AI systems are unexplainable. It depends entirely on the specific architecture. If explainability is a key part of the solution you want, you need to take those architectural concerns into account in the design of a system right from the start.

Q191 **Chair:** Thank you. Just on that point, you started your evidence by saying that the purpose for the deployment of AI was that things were beyond the computational or cognitive power of individuals. You have just said that you want everything to be explainable and you want people to see those connections. Surely there's a tension there?

Professor Hopkins: No, there's not; it's about the size of a space. For example, there are potentially billions of decision trees one could create in a generative system to understand where we could go. Those are whittled down to find the solution that we now prioritise, or the molecule we now want to make that has been generated. Understanding about where individual predictions come from—why does it score highly in individual models—can then be traced back.

It is the case that the sheer scale of the approach—the decisions and the amount of data we are ingesting—is beyond human cognition. But once we find the solution, tracing it back is a much smaller problem to understand. You can then say, "Ah, yes, I can see that insight; I can see that is connected." But that is effectively the needle in the haystack among everything else that you could have got lost in.

Q192 **Chair:** Is it not the case that with some medicines we actually do not know why or how they work? We just know that they do, and they proceed without an explanation, having been tested for safety. You are setting a higher threshold than is actually the case for many drugs that we have used for a long time.

Professor Hopkins: In fact, what I described earlier—the EXALT-1 trial—embraces what you just said. It is saying, "What we are looking for at the cellular level is the effect". It is not saying how it works. We just want to say, "Does this drug work for you?" That is the question we are asking. We are not yet asking why it works for you. We can start to work that out by integrating all the other types of omics data. That sort of backpropagation to understand that is exactly what we are doing to find new drug targets and to find out whether there is a new drug target we could go after from that result, from your test.

Q193 **Chair:** It is not a requirement to explain it as so much as it is to isolate—

Professor Hopkins: No, it is not a requirement. For a drug to be approved, we do not necessarily need to know how it works. What we need to show is the evidence that it works.

Q194 **Aaron Bell:** You have said, "we need to change the way drugs are invented". Have you got there yet, or are you still on that journey?

Professor Hopkins: It is a continuous journey because you never want to stop trying to improve your processes. We now have the first examples where the use of these machine learning techniques with our expert drug hunters are starting to show that we have produced precision-engineered molecules in a way that is far more efficient than traditionally has been the case. We will continue to do that. We continue to develop those systems. Exscientia is now asking how we scale this. How do we take what is potentially at the vanguard of this sort of technology and ask: how would

all drugs be invented in this way? We are in the period of scaling the company and scaling the technology to apply to many more projects.

Q195 **Aaron Bell:** You described earlier how you were looking at a tenth of the number of compounds as before—250 to 2,500. How much have you learned from Google's AlphaFold and what it did with proteins? Is that relevant to what you are doing? Is yours more about chemical compounds rather than proteins? I appreciate they are all chemicals, but have you absorbed that learning into what you do?

Professor Hopkins: Yes, it is part of what we do. We think of it as an end-to-end process, as a particular solution to one important problem. We have already been incorporating that technology into our early stage approaches, actually, but we are also taking some of the learnings and applying them to other technologies and techniques. For example, how could we take that kind of artificial intelligence framework and apply and develop it to be faster, against things like how you design antibody drugs using that kind of AI approach? That is where we are now working as a company. We want to think now about how we could design antibodies using AI, as well as designing small molecules.

Q196 **Aaron Bell:** Do you think we will eventually have a drug designed entirely by AI, or will AI always be there to shrink the space and to then hand it over to the more traditional approach? Or, do you think we could get to a situation where literally the whole thing is handled by AI from start to finish?

Professor Hopkins: Eventually I believe that is theoretically possible but, at the moment, the interplay—this is important for all the industries developing AI—between the human and the system is incredibly vital. A good example of this is the observations of Garry Kasparov in how AI has been used in chess over the past nearly two decades. It is called centaur chess. He found that systems could make the average chess player much more proficient and games could become tactically flawless. You could have an average player or master start to out-compete a grand master, even if you are both using algorithms, because they use different applications and systems. The real advantage comes from the combination of human skills and human creativity, with the compute power of AI, integrated into new processes. It is humans, plus machine, plus process. That is the key thing you need to think about if you want to revolutionise the industry.

Q197 **Aaron Bell:** You said it would be theoretically possible, and it might happen, that it would be purely AI. What sort of date would we be talking about if it were possible? Is it in the next ten years, the next 20 or the next 50?

Professor Hopkins: In terms of drugs that are moving forward to a clinic that have had generative AI systems involved in a design very closely, that is already happening today. Of course, we still have human experts guiding and driving the process. Human drug hunters know what they want. They know what good looks like. I do not want to put a date on it yet, but it is an active mission of Exscientia. Our mission is to solve drug

discovery. I believe that if you follow this logically through and follow the strategies through, it is entirely possible to do that.

Q198 **Aaron Bell:** Frank Nestle, the chief scientific officer at Sanofi, who you will obviously know has said that the use of AI in medical research is reliant on computer power, algorithms and data, and that the "biggest need for improvement" relates to data. Do you agree with that? What are the shortcomings of the data at the moment and how are they being addressed?

Professor Hopkins: There is still a lot to be done in terms of algorithms, particularly referring to your last question. As we move up the levels of thinking, particularly about how you incorporate human creativity and human strategy into processes, that is the real value add that human drug discoveries bring.

In terms of data, 40% of the staff that we, as a company, employ are experimental biologists in the laboratory, generating data. You need to have high-quality data to generate your algorithms. That's why large data generation initiatives, such as UK Biobank and Genomics England, which I mentioned earlier, are incredibly important national assets in this space.

Where we are focusing our data generation is working with hospitals and clinicians, and gathering those cancer samples, because the wealth of information that can then be discovered within that is incredible. Every time we invent a new technology, or a new kind of omics, it allows us to have a new insight into our biology, but that omics by itself is not enough. You want to integrate it with all the other data you can generate. That is why we, as a company, need the ability to consider how you can access patient samples and then develop machine learning, deep learning driven analysis of the samples. I believe that is where the key advances will come in personalising medicine.

Q199 **Aaron Bell:** You just talked about the importance of wide repositories of data that your firm and others can rely on. Presumably you generate a lot of your own data while doing the work. Is that all proprietary or do you have mechanisms for sharing that and making it openly available?

Professor Hopkins: No. For that proprietary information and for many of our arrangements, we have ethical frameworks within which we work with various hospitals. That is as it has to be, actually. There are potential opportunities when thinking about how to gather that data on a national basis. Again, that could be an important role for the NHS. There is data beyond even the electronic health records that were talked about earlier. In terms of how we think about pathology, biopsies and so on—it could be digitised and start to be provided into a framework. There is a huge amount of wealth, to the benefit of individual patients, that could be generated.

Q200 **Aaron Bell:** As we discussed in a previous inquiry, the NHS is a huge resource for firms like yours. What is the fair price for you to pay to the NHS for the data that they have?

Professor Hopkins: I can't really give a comment on that. We will have to sit down and work it out, as we do. For example, yesterday we announced a research collaboration with one of the biggest cancer hospitals in Germany, called Charité. That is a two-way relationship. It's not about money. It's about how you bring the data and the technologies back to the patients. That's what physicians really want. They want the data to understand how to treat their patients, so we see it as a two-way relationship.

Q201 **Carol Monaghan:** Thank you very much, Professor Hopkins. It is very interesting. Who are your current customers?

Professor Hopkins: Our current customers are everyone from large pharma companies, such as Sanofi and BMS, to small biotechs, academic groups and, of course, ourselves.

Q202 **Carol Monaghan:** When you're doing this research, you've talked a lot about the ability or the capacity of AI to look at all the different options; I think you used the phrase "needle in a haystack". Would you consider AI as an augmentation tool or will it ultimately remove humans from drug research?

Professor Hopkins: Right now it is a tool used by humans. Humans really understand the strategy we want to follow and, as I said, know what good looks like in terms of developing a drug. We are thinking about and exploring how to bring a much more strategic level into the long-term development of these technologies. It's not about replacing humans. It's about how to make people more productive.

Q203 **Carol Monaghan:** So it's not about replacing humans, but if your company does drug research so much better than a whole pile of humans working in the same field, surely it will ultimately replace humans?

Professor Hopkins: If we ever got to that point, the other benefit to society, of course, would be potentially many more medicines being available to people. It would potentially be that we then started to create a healthier workforce. We then have other economic benefits to the patient and society, and potentially the impact on healthcare costs. I think that the long-term benefits of changing an inefficient and costly approach to creating vital medicines will benefit society in the long term.

Q204 **Carol Monaghan:** What are the risks to replacing the inefficient and costly humans?

Professor Hopkins: It is not about replacing humans; it is about how you make your drug discovery scientists more productive. For example, you could be a drug discovery scientist in a large pharmaceutical company, and you could spend your entire career working on projects and never see any of those ever become a drug that goes to market. Even if you get a drug to the clinic, 19 out of 20 will still fail. Those failures also mean that patients are being exposed to those tests. Their hopes are being brought up. They are being exposed to new medicines, and you do not know all the risks that they are exposed to, and so on.

Q205 **Carol Monaghan:** That would still be the case with drugs that you would be able to produce.

Professor Hopkins: Potentially, but if we also think about how the personalised medicine and precision medicine then starts to develop, you will also hopefully start to see increased response rates in trials as well. I consider that to be a waste. What we want, and what I am pleased at when I see our drug discovery scientists working at Exscientia, is that within a short period of time people in their careers are now seeing drugs come into the clinic, moving forward and so on. That is an important element as well.

It is about how we ultimately take all the investment that we have made into education and science, and increase our ability to translate the billions of pounds of investment that we make in basic research into more new medicines for patients. The NIH did a study of how basic research funded by the NIH was translated into new medicines. It can sometimes take up to 25 or 40 years from the publication of that research until it becomes of benefit to society as a new medicine. Ultimately, what we want to do with these systems is improve the odds of doing that and decrease the time.

Q206 **Carol Monaghan:** I get all that. I am trying to understand where the human oversight is and where a human working for your company is. What is their role now? They are not going to be in the lab doing all these tests.

Professor Hopkins: Today, as it stands, they are running the projects, making the final decisions on everything and defining the strategy. It is still very much a human-driven process, with enhancements by new computational techniques to try to improve efficiencies. What you are describing is some future world, which potentially we may end up in, but I think that if we ever get to that point, the other benefits to society would already have been proving themselves in that industry.

Q207 **Carol Monaghan:** To get to that future world, Government have to play a role in regulation. What would be the key considerations for Government if we are looking at the regulation of AI in drug development, for example?

Professor Hopkins: We already have multiple layers of regulation that have been well thought through, whether they are the safety regulations that we have for medicines, the ethical regulations that we have for clinical trials, or the privacy and confidentiality regulations that we have on data governance. All these come into play.

Q208 **Carol Monaghan:** So if we are looking at AI in medical research, nothing new is required.

Professor Hopkins: No, I think the risk is potentially in over-regulating a new industry before we have worked out how we are going to apply AI to different processes.

Q209 **Carol Monaghan:** I think the fact that you are talking about over-regulation would cause some nervousness for a lot of people. There has

to be public confidence and buy-in for many people, and knowing that this had some sort of oversight would give people more confidence and more assurance that the process was going to work in their best interest. You have said that Government should not over-regulate, but there must be some areas that should be considered.

Professor Hopkins: In areas where people's decisions are being made automatically, particularly if that would affect someone's outcomes, then absolutely. There are potentially frameworks to understand how those decisions are being made—are they being made automatically, or are they being made—

Q210 **Carol Monaghan:** Would that mean needing to know how the AI tool was operating?

Professor Hopkins: No, it doesn't need to. I think that's more of a case of oversight—that is where the question lies. Take, for example, the EXALT-1 trial that I described earlier, which recommends which potential medicine a patient should take but does not actually prescribe. The actual prescription decision is still made by the physician looking at all the data that is been generated and saying, "This is why, looking at this data, we believe that this medicine might give a better response than this other medicine." It is fundamentally the physician's decision to make. They are responsible for the oversight. The systems we are talking about here are ultimately about how we bring better information and recommendations based on that information to the responsible party. However, it is still in that responsible party's interest for it to be responsible for that decision, as it is today. It is fundamentally about how we provide new information.

In that sense, I do not see AI as a revolutionary force. I think it fits within the exact societal frameworks we have today. It is just a new approach to how we can potentially be more efficient in the use of data. That's why I said that the risk is actually not applying AI in many cases, because the risk is business as usual, which in many cases means inefficiency and cost—to not just the healthcare provider but the patient. We need to look at the possible advantages as well, and understand that if we don't embrace the possibility now of better decision making and better information, we are making a choice to lock ourselves into today's framework. If we are happy with the way healthcare systems run today, that is fine. I'm not.

Q211 **Carol Monaghan:** Okay. Changing tack slightly, you are obviously running a successful company using AI in the research. What are the advantages of the UK regulatory environment over other jurisdictions, and are there any less advantageous aspects of the UK's environment at the moment?

Professor Hopkins: That's a really good question. I think the UK has actually led the world in thinking about some sensible frameworks, as I said, within data governance, safety of medicines, ethics and so on. It's a well-thought-out approach. The UK is also a leader in how we think about the use of AI in healthcare and AI generally, actually. We have been, to

use a phrase, punching above our weight in the early development of this new technology.

We are not necessarily putting as much investment from both the private sector and the public sector into where this potentially could go. However, it is interesting, if you look at ourselves, BenevolentAI and others, that it was UK companies who were some of the first into this whole area. I think that reflects the environment the UK has had in terms of both the academic backgrounds of producing investment into AI and the data backgrounds, as we highlighted earlier. The UK has been fantastic at creating large data assets. All of that creates the right environment and the right ecosystem for these initiatives to flourish.

What the UK needs to think about is that we have a very strong life sciences sector at the moment, in terms of public investment and early stage biotechs. Because of certain pressures right now, in terms of the cost of drugs—the pharmaceuticals market—the UK is probably less attractive to many international companies from that point of view. We don't necessarily play to being an attractive market from the point of selling drugs. However, we can be an attractive market as a place to do research. That, I think, is where the UK's strengths could be. That is why having the right governance framework is important to encourage research and new ways of how things could be done, and also to think about how large data assets, as we mentioned earlier, could be an important part of attracting research to us.

It is beyond this Committee, and it is beyond me to comment, but there's lots to do with how we think about drug pricing and things like that. I am not going to talk about that, but I do think we can control, and we are in a very good position, on the supply side of things, how we supply researchers for our education system and how we supply for the possibility of large data assets for academic—

Q212 **Carol Monaghan:** But you have said that funding is problematic or challenging.

Professor Hopkins: That is not necessarily about Government funding—that is much more of a comment.

Q213 **Carol Monaghan:** Where else would fund research better? Are we talking the US, Germany—those kinds of places?

Professor Hopkins: We listed on the Nasdaq in October '21. The reason we listed on the Nasdaq rather than on the London Stock Exchange is because fundamentally capital is much more fluid to attract in the US.

We have an incredibly international investor base, from Novo in Denmark to SoftBank in Japan to Baillie Gifford in Scotland, to many large US funds as well. What we found is that the stories we can create in the UK are attractive internationally. What was hard is that we built a company, I would say, despite being in the UK—it wasn't the support of UK investors who built Exscientia; it was the support of international investors.

The issue is not necessarily around Government financing in the UK or Government funding. The bigger issue for developing AI as an industry is actually that point between the start-up and the IPO and about how one attracts international capital into the UK market.

Q214 **Carol Monaghan:** Thank you. My final question is again about regulations. You talked about the requirement to have decent regulations; I cannot quite remember your exact words. Do we have sufficient capacity to develop proper regulations? Is there sufficient expertise in this area?

Professor Hopkins: That is a really good question. You do need to understand the individual domains. AI being applied in each different industry is going to require a different framework.

The way to approach it more generally is to think about what the key societal issues one wants to address are—issues of privacy, an issue potentially around decision making and agency and autonomy being taken away—and who is responsible for decisions. I think probably that is a better position to start from, because individual regulation is going to be down to each of the individual specific industries that are being represented and are probably presenting evidence to you, and those are very specific. But I think there are some more general worries and I think they are probably sitting closer to the area of how we thought about data governance in terms of privacy and confidentiality, and the governance frameworks we put in place for data.

Chair: Thank you very much indeed. I will go to Stephen Metcalfe, finally. If we could keep things as concise as possible, we will then get on to the next panel.

Q215 **Stephen Metcalfe:** I will do my best, Chair. Following on from that, are there any international examples that we should either look at to improve our system, or avoid to make sure that we don't fall into the same traps?

Professor Hopkins: Good question. I think this field is still so early that no one yet has made two bad decisions or two good decisions. That's the important thing to remember—we are still innovating. We are still evolving. This field is only just now breaking open.

We need to understand how we can apply AI in lots of different kinds of industries. I think the key advantages—and this is what was interesting about the national strategy that was proposed last year—is that it really is fundamentally about how we can help many companies start to adopt AI. I fundamentally believe that the companies of the future will be AI-first companies, and those that don't think like that will be left behind.

Q216 **Stephen Metcalfe:** Okay. If no one has quite got it completely wrong or right yet, do you think there should be a harmonisation or co-ordination of regulation globally? Or is there a competitive edge to be had from the diversification of regulation?

Professor Hopkins: You could argue for both, but considering that, for example, many systems are obviously transnational by definition—the

internet is—then if there are particular issues around privacy or around agency, particularly probably in more of the social, general and public areas, then there is actually a strong argument for thinking about what the approach would be that one would take on a multinational level. You can use a website from anywhere in the world, for example.

In terms of thinking about specific industry applications, where you want to make industry more dynamic and more efficient, there is potentially a competitive advantage there, and also a potential connection to other investments and other strategies that a national Government can undertake, such as, as we talked about earlier, the investments the UK has made into large data generation, and the potential assets we have with NHS data and the patient material there.

Q217 **Stephen Metcalfe:** Perfect; thank you. I have a final question. The Government are about to publish a White Paper on AI. Is there anything you would like to see in it? This is your chance to bid. Have you as a sector been consulted on it?

Professor Hopkins: I don't think we have been consulted as much as we should have been. It is wonderful to present here today, but we have not had as deep discussions as I believe we could have had. I believe the UK can see a particular advantage through the interaction of AI in healthcare, and I would like to see within that paper an approach that thinks about elements of data generation as a key part of that advantage. Ultimately, AI is driven by quality data, so we need to think about our data assets as a key part of our strategy.

Q218 **Stephen Metcalfe:** I said that was my last question, but can I clarify something? Am I right that you are saying that as we move forward, if data is going to be generated or gathered as part of any process, we need to think about how we might make that data accessible to the AI sector more widely very early in that process, to build that in at an early stage?

Professor Hopkins: Absolutely; you need to think up-front about confidentiality and ethics issues right from the start.

Stephen Metcalfe: As opposed to trying to do that backwards with existing data.

Professor Hopkins: Absolutely.

Q219 **Dawn Butler:** I have a quick question. If I had a company and I used quantum computing, would I put your company out of business?

Professor Hopkins: No.

Dawn Butler: Why not?

Professor Hopkins: First, we are developing quantum techniques alongside AI. We believe that the future is the integration of AI with physics-based techniques. Secondly, what is interesting about creative AI is how one really thinks about the interaction between the human

elements—understanding human creativity—and how one ultimately encodes that. This is beyond just a simple algorithm. It is actually about a whole ecosystem of algorithms and how they interact. The mission of Exscientia is to encode and automate drug discovery, and that means ultimately understanding the human creative element of it as well.

Q220 **Chair:** Among many interesting things you have said, Professor Hopkins, one was that people in the research side of the pharmaceutical industry could go through their whole career without having a drug that actually works with patients. We have heard evidence to this inquiry that 86% of drug development programmes between 2000 and 2015 were unsuccessful. You had a long career with Pfizer. Did you find that the fact that people were literally committing their life to something that didn't come up with anything beneficial was becoming a problem in terms of attracting people to make careers in pharmaceuticals?

Professor Hopkins: What is fascinating is that the people I know who work in the biotech and pharma industry, and in the life science academic world, really are committed to this space. Many could probably earn higher salaries elsewhere, but they are really driven by a mission. People are really dedicated: they want to see their skills, their knowledge, and the science they do really apply to the creation of new medicines. It is very much a mission-driven vocation.

Q221 **Chair:** But outside that, did you experience that people were being put off of going into the sector once it entered a phase of much less certain success?

Professor Hopkins: People were potentially being put off by the lack of success, as you call it, in that the economic implications, as you might remember, led to a very significant number of job losses in the UK pharmaceutical industry. Many large research sites closed down and many companies moved back to the US. The impact has been a reduction in the pharmaceutical industry in the UK over the past 20 years. That is fundamentally what I would like to change.

Q222 **Chair:** So if AI is successful in drug discovery, its economic impacts could be transformational.

Professor Hopkins: I fundamentally believe this is the next phase for the pharmaceutical industry.

Chair: Thank you very much indeed. You have been very generous with your time and we have gleaned a lot of very useful information from you.

Examination of witnesses

Witnesses: Professor Delmiro Fernandez-Reyes and Professor Mihaela van der Schaar.

Q223 **Chair:** As our next panel of witnesses take their seats, I will introduce them. Professor Delmiro Fernandez-Reyes is professor of biomedical computing at University College London and adjunct professor in paediatrics at the College of Medicine of the University of Ibadan in

Nigeria. Professor Fernandez-Reyes is also a medical doctor and has pioneered the use of machine-learning approaches to further understanding childhood malaria, among other things.

We welcome, alongside him, Professor Mihaela van der Schaar, who is professor of machine learning, artificial intelligence and medicine and director of the Cambridge Centre for AI in Medicine at Cambridge University. Professor van der Schaar's main research areas are the application of AI and machine learning to medical and healthcare challenges—that is obviously very relevant to our inquiry—and she is a Turing faculty fellow at the Turing Institute; so many affiliations.

Thank you both very much indeed for coming and giving evidence today. Perhaps I could start with Professor van der Schaar. You were involved in the Topol review, which was commissioned in 2017 by our colleague Jeremy Hunt in his then capacity as the Secretary of State for Health and Social Care. The review published a report in 2019, looking at technologies in health and medicine. From 2017, a panel that you co-chaired looked into the impact of AI. Could you reflect on what you thought then and how things have been since then? As we know, you reported in 2019, but three and a bit years is a long time in the development of AI.

Professor van der Schaar: Thank you very much indeed. I did co-chair that panel. For context, the Topol review had three different strands: one was genomics, one was digital medicine, and the one that I co-chaired was indeed related to AI. I would say that the pandemic somewhat indirectly helped in the digital medicine aspect of the panel, because all the clinicians were using technologies to serve their patients at a much faster pace than anticipated.

In AI, I would say that some progress has been made. Last year, Health Education England established 10 fellowships for clinicians who are, part time, learning and developing technologies in AI, and 10 new clinicians will be enrolled in the programme this year. I would say that 20 clinicians so far is a relatively modest number, although in this case, of course, the pandemic has hurt.

What I would have hoped to see is much more involvement with AI in the clinical space. I personally do not want to turn clinicians into machine learners, but I really would have hoped that they would understand these technologies a lot better, not only to adopt them but to shape them. What I was really very passionate about in that particular report was for clinicians not only to be at the receiving end of these technologies but to be much more involved in developing them and providing the right way of thinking about what we really need.

Q224 **Chair:** Thank you. Was it your thought, and perhaps is it your thought, that AI would go from being a narrow and rather specialised field to being mainstreamed—the word used in the past—or spread across the sector? Or would you expect an expertise to be retained within specialists but that they would become more influential as an expert group?

Professor van der Schaar: First, I see AI purely as a tool, not only in medicine but in transportation, finance, businesses and many other domains. I worked in many areas prior to falling in love with the area of medicine. I think this technology can really empower many domains, but I do think that currently too much is pushed by people like me, in machine learning, imposing these technologies on others, rather than educating them about them and allowing them to develop these technologies.

I see AI more or less as a tool that, for instance, in this case clinicians could shape so that they develop tools that are useful for them. In that regard, it is important that they are aware of the possibilities as well as the dangers of the technology, so that the right tools are built rather than pushed on to them.

Q225 **Chair:** Finally from me, before I turn to Christian Wakeford, your strand of the Topol review, reporting in 2019, said: "AI will transform patient-generated data into clinically useful information and empower patients to manage their own health or seek appropriate health support. Patient benefit should be the driving force behind AI…with new products co-developed with patients from design to implementation." Is that still your view, and will you comment on how it has happened in practice since you wrote those words?

Professor van der Schaar: It is very much still my view. As a matter of fact, while I do think that drug discovery is a very important aspect of our healthcare, I think it is even better if we are never getting sick, if we can prevent disease, and if we can identify disease early so that maybe we don't need to be treated or we need to be treated minimally. I think, as a potential patient, that this is very important. I think it should be at the forefront of any technology in AI. In my opinion, too much is focused on treating disease rather than using AI to prevent disease, to make sure that we identify early disease and, if a patient develops a particular disease, gets the right intervention at the right time in the right location and according to their wishes.

Q226 **Chair:** Thank you very much. Just before I turn to Christian, I noticed that Professor Fernandez-Reyes was nodding at what you said. Perhaps you might add a comment, Professor Fernandez-Reyes.

Professor Fernandez-Reyes: Yes, I certainly agree. A focus towards preventive medicine is important, because preventive medicine certainly reduces cost significantly across public health. You were talking about the training of clinicians in the domain, which is very important and has been happening over the 20 or 20-something years in which the MRC has been focusing on that. In fact, I am a product of that, because I am a trained computer scientist too. I specialise in machine learning—so I am not just a medic—in that area. There is a two-way avenue: it is not only that we need to train clinicians to understand; we also need to take the machine learning, the people who do the technology, into what are the domain's specific characteristics.

Q227 **Christian Wakeford:** Professor van der Schaar, you have suggested two

distinct AI categories: "petri-dish" and "reality-centric". Looking at both of them, it would appear that reality-centric is more suited to medical research. Can you explain why? Can you think of any circumstances in which the petri-dish approach would be more appropriate?

Chair: As you do so, will you just explain them, for people tuning in who may be unfamiliar with those terms?

Professor van der Schaar: As a matter of fact, it was relatively recently that I coined these two terms, partly because of concern about the way in which the machine learning and AI community to which I pertain is going. I see reality-centric AI as AI that can operate effectively, accountably and reliably in the real world and which focuses on problems that are human-centric and messy. I liked it when a previous member of your panel referred to humans being imperfect and messy. It is very much focused on that. This is, I would say, the core of AI for healthcare, AI for education and AI for business.

That is opposed to what I call petri-dish AI, which is very much focused on developing AI technology for playing games or for solving technological problems or scientific problems that are very well defined—for instance, protein folding. This is wonderful and I think it needs to continue. But really, the world is not a board game—a game of chess or even of Go. The reality is, in these settings, that we operate under complete information; everything is known. The rules of the games are predefined. But the reality of the world—if we take something like the NHS, it is complex, messy data. I would say often the environment is changing according to rules we cannot really predict. Events such as pandemics happen. Data is not, "The more, the better." I can play billions of games of Go, and learn the game. I cannot reorganise the NHS billions of times. I have written this piece to concentrate the minds of my community—the artificial intelligence community—but more generally the world to focus not only on the hyped area of solving games, but on solving problems in the real world, such as problems that the NHS encounters as a result of limited resources or complexities such as this pandemic.

Q228 **Chair:** To be clear, a Petri dish is a protected artificial environment, as opposed to the real, messy world.

Professor van der Schaar: And it is very well defined and very clear. For instance, determining proteins and protein folding is a complex and important but very well-defined problem. If I am to develop AI to improve the performance of the NHS and to empower clinicians in the NHS, that is a much less clearly defined problem, but it is even more important and exciting. I would like my community to focus much more on real-world problems and not just on very niche areas.

Q229 **Chair:** Before I hand back to Christian, Professor Fernandez-Reyes was very keen to come in on that.

Professor Fernandez-Reyes: I disagree with some of the comments that my colleague has made. Let us focus on uncertainty. Uncertainty is very linked to governance, because one reason why you want to govern is

to create trust in a technique or an approach. With machine learning and artificial intelligence, academia and other researchers in industry have tackled precisely that with probabilistic approaches, which are very successful at tackling uncertainty in a system.

Despite the fact that they have done that because of games—the famous AlphaGo of Google and others—the important thing is that with those frameworks, we can tackle uncertainty precisely. That field is important because of that. Now we need to extrapolate and put into a governance framework the characteristics of these systems. It is not trivial. The things we do with protein folding involve a very complex system, and it is very uncertain. I would not call it a Petri dish. This gives humans a tool that can augment their ability to tackle uncertainty. Now we need to have governance over it for our prosperity in healthcare and other areas.

Professor van der Schaar: Please allow me to intervene for a second to say that it is not just about uncertainty; I am not saying that these problems are not difficult to solve. You have a large state space that you need to solve when you solve protein folding or you are playing Go. What I am saying is that it is not just about uncertainty; it is about modelling very complex, messy systems that have humans in them. Simple games of Go or even complex problems like protein folding do not capture the complexity of the real world like we have in transportation systems and the NHS. It is to that point that I made this point of difference.

Chair: Can I ask both of you to raise your voices a little? Although you have got microphones in front of you, they are for broadcasting rather than in the room. It is a long way between you and us.

Q230 **Christian Wakeford:** This question is also to Professor van der Schaar. The centre aims to advance regulation, safety, privacy and ethical standards within AI and machine learning. How is progress so far, and why is it so important?

Professor van der Schaar: It is extremely important to make sure—especially for AI in healthcare, for criminal justice and for education—that we build very careful, transparent and ethical AI. In that regard, the Committee discussed some very important questions earlier. We cannot build high quality AI unless we have high quality data, and, unfortunately, high-quality data requires dealing with a lot of biases, a lot of noise and a lot of errors. Again, it is not just about having a lot of data. We need to have the right data otherwise all sorts of issues and biases are going to be propagated. In my opinion, not enough energy is spent on developing technology—exclusively AI technology—to improve the quality, the representation and the biases in the data. That is something that is at the forefront of my research, but I wish more people would worry about not only building analytics but making sure that the data on which they build the analytics is of the right quality.

At the level of algorithms, it is, again, very important to make sure that the algorithms work well not only at the population level but in minority

groups, and that we audit the algorithms not only for global performance but for individual performance.

Q231 **Christian Wakeford:** It is like you're saying the data is messy. Professor Fernandez-Reyes, considering your background as both a biomedical scientist and a medical doctor, can you explain the benefits of AI to both fields of work?

Professor Fernandez-Reyes: Of course. I am a computer scientist, too. The benefits are many. In terms of what we were talking about before, most of these techniques are diverse in their own right. There is not just a single AI thing; it has evolved over 30 years, with three waves of it.

The current moment, and the past decade, has helped us to understand, for example, the complexity of phenomena, especially in biological systems. In my case, I focused on understanding the pathology of malaria in children—why some children develop severe anaemia and need transfusion. That is a very complex disease. It has many factors, from the population to the child's genetics and all the components of the actual infectious agent.

In a way, some part of these techniques or these paradigms allow us to take very complex data, as the previous witness said, and actually make sense of it and find patterns in the data. These patterns might be important. Once we have patterns, we can reverse back. It is a very difficult phenomena, in the sense that you go to a reductionist path. You form hypotheses. To give you an idea, when we talk about omics, we were in three dimensions with what we see, but when you start adding variables that becomes like a cube of many points that we cannot imagine. It is very complex.

These techniques have been very useful over the decades to form finely identified patterns. You can then go back to the lab, find patterns and find these famous things called biomarkers. How can I simplify this in a biomarker that allows us to stratify risk? For example, what the previous witness was saying is precisely how we stratify patients according to whether they will benefit from a drug or the risk. That is the base model part of it.

The second part is how do I deal with that? What do I do with it? Diagnosis is a very complex matter. A human—a medic—trains for 10 years minimum, and is basically a complex machine associating many things, but he suffers from cognitive tiredness, cognitive load errors and biases, because it is impossible not to have biases. In terms of diagnosis, this can augment the work of a clinician or a clinical care team—because it is not only the clinician but the care team—and the augmentation that provides can speed up referrals or, as we were saying, prevent the disease from getting worse, etc. This can be another area in which there are clear advantages that we can put in.

Of course, there are other areas that cannot translate into a treatment of disease. In my case, we look for explaining the disease, finding

biomarkers and using that as a diagnosis preventive. Where can I find those individuals at risk? Of course, that is always context dependent in terms of the disease. The biomedical domain will be like that. Too often, the conversation we have focuses on diseases that are characteristic of long-living individuals in a population—such as cancer and neurodegenerative diseases—but very little attention is paid to, for example, infectious diseases. It is a completely different approach. Infectious disease you need to treat quickly; otherwise, you will see the effect quickly.

The use of that is how you govern this space. That is what we are doing here, in a way. It is true that I can develop all those things, but how can I put in place the systems that the person who is going to benefit the most, which is the clinician, trusts so that he is going to take a clinical decision that will not legally affect that person and the patient?

Christian Wakeford: Thank you. I think that actually answers the next question as well.

Chair: Thank you very much indeed. That was a great question. I think Katherine wants to come in on that.

Q232 **Katherine Fletcher:** I promise I will be brief. I want to make sure I understand what you are saying. You are talking about complexity in systems. In human and biological systems, that is effectively a bell curve distribution. If I understand you correctly, you are saying you need data that is representative of the entirety of the bell curve—for proteins on the outside of the coat of the Plasmodium falciparum, MHC complexes within the individual child that is getting malaria or something like that. Are you saying that, for this to be effective, you have to give clinicians confidence that the overlay of all those different distributions of bell curves is within the data, to allow the AI to not spit rubbish out?

Professor Fernandez-Reyes: Not necessarily. As I say, it is context-dependent. It is not necessarily bell-shaped. You are talking about the distribution of the individuals that generate the states. The data can be multi-dimensional; it can be millions of data points. You can choose the characteristics of the data point; it might be bell shaped or not. In a way, that is what these AI techniques are very good at—doing complex models to solve discriminatory patterns. However, it is very difficult to explain, precisely because, as I say, it is not as simple as a Gaussian shape.

But coming back to the answer to that, I do not build the same AI system, for example, to detect malaria in, let's say, Ibadan, where we have a high prevalence of the disease—it is very frequent and very common—as I do to put in Heathrow to detect which travellers have malaria. It is completely different.

So you can get techniques to measure these things, but when the MHRA started looking at this in 2022, one of the high-level topics it highlighted was that there is no framework in which—

Q233 **Katherine Fletcher:** Yes, that is what I am trying to get at. We are

trying to make sure that this stuff is understood and that it does not produce stuff that potentially lacks a benefit. You are saying that there is lots of potential, but the data is important. In terms of protein folding and quantum effects, how on earth do you put that into—

Professor Fernandez-Reyes: But it is context-dependent as well.

Katherine Fletcher: Okay. Thank you. I am conscious that I am stealing Aaron's time.

Chair: We are going to come back to you, Katherine. We will go to Aaron, then Katherine, then Graham and then Dawn.

Q234 **Aaron Bell:** Thank you, Chair, and thank you, Katherine—*[Interruption.]* No, not at all—it was a great question. We have just talked about opportunities. I want to ask you both about the risks. I think Professor van der Schaar talked about them a little already, but from your perspective, Professor Fernandez-Reyes, what are the biggest risks of using AI in medical research?

Professor Fernandez-Reyes: I think the biggest risk we have is, first, we don't know what we don't know in the implementation. Of course, there are the risks we have discussed in this session—the risk of not being transparent and the risk of giving an explanation that may work, but we don't know how to do it.

Safety in the domain is critical. Are we creating problems? Are we creating systems that are not safe or, therefore, robust? That is one of the key issues at the moment—a key area of research—in biomedical and medical research and AI systems, which are deep, deep representations. I don't know about protein folding, and all that; I have done deep models or deep machine learning.

Those models are actually very susceptible to small changes and can produce a completely different outcome. That is one of the risks we need to analyse. As the medicines regulatory agency was saying, that boils down to us needing to create datasets and tools, so that I can create guidelines about how to test that this system is safe and robust. That is not trivial; it requires research in its own right. Creating that toolbox requires an interdisciplinary team and stakeholders, which could include data scientists, mathematicians and other people, including clinicians and the regulatory agency.

I was happy when I read the October 2020 report. The MHRA has three work packages precisely on AI as a medical device. That is the beginning of the conversation. I think there need to be resources for that. The safety framework of medical biomedicine is there; now we need to put this AI, as a medical device or a medical framework, on top of that. That relates precisely to these issues of lack of safety, transparency, interpretability, representation and inclusion, as Mihaela said before. Those issues need a framework to show people how to comply.

Q235 **Aaron Bell:** Thank you. I will come to you in a second, Professor van der

Schaar. Are these new risks, or just evolutions of the risks that have always been there in medical research?

Professor Fernandez-Reyes: With safety, there is always a risk.

Q236 **Aaron Bell:** Yes, but also inclusion, bias and so on.

Professor Fernandez-Reyes: The risk has always been there. I give you a very simple example. When children develop, we measure them with a curve. This curve has developed mostly for people living in the north. Over the years, my colleagues at the University of Ibadan and I have had to create one for sub-Saharan Africa—this is just a table on a piece of paper; I am not talking about AI.

These problems exist—problems of sampling populations and inclusion. This is on a larger scale now, of course; with the internet and globalisation the genie is out of the bottle, in a way. The problems are there, but they are exponentially bigger.

Q237 **Aaron Bell:** Thank you. Professor van der Schaar?

Professor van der Schaar: I briefly want to reiterate that explainability, transparency, bias and trustworthiness are all very important, and I work a lot on them myself. But it is not said enough that medicine and healthcare are human-centric subjects.

One of the dangers is that if we do not think carefully and regulate technology carefully, we will build methods that consider the patient as a square. We will make the wrong abstractions about what patients need, and about the complexity of patients and healthcare. That is why I am going back to this reality-centric agenda: I really want to see the human—the patient, the clinician—at the forefront.

One of the dangers of building AI that aims to be artificial general intelligence is that we forget about humans. I liked very much the comment made earlier that humans need to be at the forefront—this is about them and their needs as a population and as individuals. I am concerned to some extent about the lack of emphasis on the human-centric aspects of this technology.

Q238 **Aaron Bell:** How should we address that, as policymakers?

Professor van der Schaar: What is really important is to make sure that we take a human-centric view rather than "profits first" or "drugs and innovation first". For instance, let me backtrack a little. There is a lot of discussion about cutting-edge technologies for drug discovery, and those are very important. But machine learning can play an even bigger role in doing the dirty work—improving logistics in the NHS.

During the pandemic, my group—just a professor and two students—joined forces with NHS Digital using, I think, the first machine learning tools used in the entire UK, to predict how many beds and ventilators we would need and who would need them.

We talk often about this lofty AI and how we are going to know everything about the human, which is wonderful, but we can use AI for logistics to solve problems of the type of resources we need and when, and to predict when a particular patient should return. We talk about giving people drugs, but the reality is that at times a patient may be at very high risk, and currently AI is used very much on the basis of one-size-fits-all.

You can identify which patients should come to be screened and how often; if they have disease, when they should come back and how often; and the associated competing risks. It is important for us to consider these complex and messy issues, but we forget about them.

Q239 **Aaron Bell:** So when it comes to improving the health of the nation, you would say that we should focus much more on systems and diagnosis?

Professor van der Schaar: And the logistics part of it.

Q240 **Aaron Bell:** Yes; put some of Amazon's AI into the NHS.

Professor van der Schaar: But with the realisation that we are not packages and the focus is not profit. The focus is really the healthcare of the population and using this technology as empowerment rather than creating incentives and markets that may not be there.

Q241 **Aaron Bell:** Do you want to comment further, Professor?

Professor Fernandez-Reyes: I agree with most of what Mihaela says. Coming back to what you were saying, we were talking about the Topol report before, but it was a subsequent report by NHSX that put all the AI and governance in there, so I think that that is worth revisiting. The Topol report didn't go into governance and focused mostly on AI and robotics. Let me come back to diagnosis, prognosis and outcomes, and prevention, as we saw before. Treatment, which is the drug discovery part, is an important axis—

Professor van der Schaar: It is just emphasised too much—

Professor Fernandez-Reyes: Exactly. What I am trying to say is that this is a whole—from when an individual is born in a society to the completion of their life, these factors act in a complex interaction. Can I ask a question? Is that possible? Do you know how many AI systems are implemented in the NHS? How many AI devices? Not a lot. Just recently, in 2022, NICE evaluated the use of AI systems for diabetic retinopathy. AI systems are very good at looking at images, so the low-hanging fruit at the moment is medical imaging. It is there in a variety of cases. Diabetic retinopathy is a low-hanging fruit that we can assess. NICE reported on that in 2022. The other important one, which is coming soon, is the automatic analysis of CT scans of individuals with suspected stroke. That has tremendous potential for us as a society.

Those systems are already being evaluated by NICE together, of course, with the regulatory agency, but we found that the framework in the regulatory agency was not there, despite our already saying in 2016 that now there needs to be regulation, but that it needs to be agile. That

comes back to what the previous witness was talking about; regulation must exist, but how do you deal with the complexities and how are you as a regulator agile enough not to stop innovation? I would say we need to put resources and knowledge in institutions like the regulatory agency—the UK regulatory agency is quite mature and good—and put in interdisciplinary expertise.

Let me give you an example—

Chair: Briefly, if you would.

Professor Fernandez-Reyes: The previous witness spoke about medicines and trials, but he did not mention the most important problem they have. Clinical trial methodology has not been developed to understand more than 10 variables at a time. With AI, it is crazy.

Katherine Fletcher: That is what I meant about the bell curves. I have a suggestion I want to test with you later.

Professor Fernandez-Reyes: Basically, clinical trial units invested heavily in this, and clinical trial units now are mature enough to bring in that AI component in what we call adaptive trials, trials that not only take place in the classical way but are adaptive. There are two new branches of doing trials, umbrella trials and basket trials. That comes to the question of personalised medicine—N=1. Those trials are the way to do it, but if we are going to capitalise on the whole thing—because a device leads to a trial and everything, even our logistic system, will need to have an implementation trial—we need to facilitate the knowledge inside our regulatory system and among stakeholders to allow agile regulation.

Q242 **Chair:** Before I turn to Katherine, we need a recognition in the regulation of clinical trials of these new opportunities. Are there other jurisdictions in the world you are aware of that are already doing this for clinical trials?

Professor Fernandez-Reyes: Remember that, as the previous witness said, the US is very big in cancer research. Yes, the US have a similar approach to us: light touch in the area of AI. In fact, they have already approved many AI as medical devices and so on. We have a light touch. The European Union is stricter and less innovative, so it has more hands-on regulation. I agree that regulation must exist in all the areas that Mihaela was talking about—transparency and so on—but it is about bringing that agility to it. The only way we do that is to train the regulator, so we need to put interdisciplinarity into the regulation, and the pandemic has shown us that very clearly.

Q243 **Chair:** Professor van der Schaar?

Professor van der Schaar: I would like to say two things. As one of the researchers in the UK, and I dare say internationally, who has done the most research on machine learning for clinical trials and changing clinical trials, I have been doing research in this area for a very long time and collaborate with people in the United States, where I was a professor for many years, as well as in Europe. Again, looking at the many steps of

clinical trials, and understanding the complexity of it, is a complex system problem, because it is about which patients you recruit. Do you recruit the right patients, or do you select only some minority groups that you are testing? Then, when the drugs end up being on the market, only certain patients may benefit. Again, it is about looking at those responses. It is not only about drugs, components and precision medicine; it is also about the right dose that is given to a particular patient. Currently, the patients included in clinical trials do not have comorbidities, and the reality is that most patients who will get the drug will have comorbidities and will be elderly.

Another very important aspect is what happens to a drug after it is approved. Is this drug going to have long-term side effects? Is it going to create additional comorbidities? It is about both regulating and improving clinical trials, but it is also about the long-term effects and how this particular drug is going to be valuable in a particular ecosystem—for instance, the NHS—that is really quite different from the US or European market. Again, there is an emphasis on the complexity of the end-to-end systems, which are very important, but I would really like to jump into the previous discussion. We hear a lot about how AI is miraculous in radiology and imaging, and I have worked in this area for many years, but there is an even better opportunity that I want to emphasise: using AI to solve logistic problems and to predict what we need, and when we need it, to be able to help with the infrastructural aspect of this.

Chair: That has come across very clearly from your evidence, and it is very useful. I will turn to Katherine; otherwise, we are going to run out of time.

Q244 **Katherine Fletcher:** I am going to translate what is very sophisticated into something simple, just to make sure. You are saying that, in clinical trials at the moment, the AI cannot have a load of blokes having HRT tested on them. If those blokes also have diabetes, the AI will come out with an answer, but it is going to be absolute nonsense. Professor van der Schaar is nodding, but Professor Fernandez-Reyes is not. I want to try to simplify it for normal human beings.

Professor Fernandez-Reyes: Comorbidities are linked to the drugs.

Katherine Fletcher: The thing is, you guys get it. I am trying to simplify it, just to give it a "for instance" that will never happen. Professor van der Schaar, your reality-centric model is effectively now splitting into two subsets as we hear your evidence. There is almost a novel discovery piece, and then there is an infrastructure, logistics and planning piece. The novel discovery, as well as the medical imaging, is getting lots of attention at the moment, and it is impressive. But the regulator has to understand what is coming out, given all the complexities we are diving into. I am going to say three things, and I want you to tell me whether they are nonsense.

Chair: You need to do it quickly, because we are running out of time.

Q245 **Katherine Fletcher:** I will. You need known results—known datasets—to go through an AI, because if you know the answer and the AI is not coming out with the answer that is known, there is a problem. You need to test those knowns going through, and then most AIs in novel discovery need to be giving probabilistic scores on results and range estimates, to allow for the uncertainty in the dataset. If a regulator did that, would it be a good place to start?

Professor Fernandez-Reyes: It is an over-simplification, but yes. Does that open a discussion with the regulator? Yes.

Katherine Fletcher: Professor van der Schaar?

Professor van der Schaar: I would say it is a good place to start, but more is needed. You really need to look again at sub-populations and make sure it is true not only at the population level, but at the sub-population level.

Q246 **Katherine Fletcher:** Yes, that would be in the known results list. Do you want me to go on to the future?

Chair: Yes, but very briefly.

Q247 **Katherine Fletcher:** We have talked a lot about the future. Let's just speculate. Do you think it is possible for a drug to be discovered entirely by AI? If it is, how far away in the future is that?

Professor van der Schaar: Again, my emphasis is not so much on just discovering drugs. It is making sure that the right drug is given at the right time, to the right patient with the right dose. There is too much focus on just discovery. There are many other complexities to make a drug effective in a population.

Professor Fernandez-Reyes: I agree that discovering a drug does not mean it is effective. It cannot be done completely by AI. The AI for doing that does not exist. We are doing narrow AI.

Q248 **Katherine Fletcher:** The AI can say, "Have a look at this", but the complexity to make sure it is safe in humans is not within an AI system.

Professor van der Schaar: And useful for humans. Not only safe, but useful.

Katherine Fletcher: Perfect. I will stop now; I got over-excited.

Chair: Let me turn to Graham. No? Well, you can carry on.

Q249 **Katherine Fletcher:** I am quite interested in chasing through this idea of screening and diagnosis, provided you test the AI with a known individual. I am a 47-year-old, middle-aged woman who likely has reduced oestrogen or whatever. You know that I am probably in that bucket. You put me through the AI and out pops, "Talk to her about HRT." Can you give us an idea of the kind of advances in screening and diagnosis? I don't think that AI exists at the moment, does it?

Professor van der Schaar: Thank you very much for that question. We have done research on numerous diseases across everything from a variety of cancers, to cystic fibrosis, to Alzheimer's, to cardiovascular disease, to endocrinology and to many others. The reality is that the technology exists today. It is a lot better than the statistical technologies—the QRISK scores—that are currently used in the UK. We have validated this in more than a million patients. The trouble is that it is not the flashy AI that gets prime-time news.

The reality is that the technology is in place. I really wish that I was not an academic. I wish I was somebody who could translate it to the NHS. But we have worked closely together with NHS Digital to build some of these technologies. I think that they could be transformative, because currently when you go to the clinician, first, you are tested for maybe one or two risks. The reality is that at the push of a button, you could be tested for many risks.

Q250 **Katherine Fletcher:** You are saying that a standard run of blood screens, with a genetic set of coding and some basic knowledge about an individual, could spit out, "Have a word with them about this."

Professor van der Schaar: Definitely, maybe even without genetics.

Professor Fernandez-Reyes: But genetics are expensive.

Professor van der Schaar: No, I disagree completely. What is very interesting here is that we think the more data, the better. We need genomics; we need that. The reality is that some of these technologies would tell us the minimum amount of information we need to know about a particular patient. We always think wrongly about AI, partially because of the wrong narrative. We think we need big data sets, genomics and expensive tests. No, we do not. For many diseases we need actually relatively little information about the patient to make a relatively good prediction.

Q251 **Katherine Fletcher:** Provided it was a probabilistic, 80% chance, and it said, "Maybe have a bit of a natter."

Professor van der Schaar: Exactly, and maybe I could increase the certainty if I did a particular additional blood test. Another important thing is that it is not only one moment in time. You have a history, so as you are evolving through your life, you may end up having an increase in a particular biomarker. You may lose weight; you may gain weight. You may have other diseases being diagnosed. What AI is good at is looking at these trajectories of disease to predict risk and competing risks over time.

Q252 **Katherine Fletcher:** Let me just simplify that again, and then I will bring in Professor Fernandez-Reyes. For example, the classic folklore would say that somebody with cancer will lose weight beforehand even if they have not presented to a doctor. You are saying that in a routine medical check-up, if there is a significant drop from a previously stable weight, that might be enough of a trigger to go for a very simple cancer blood screen. Is that kind of what you mean?

Professor van der Schaar: Yes, that is kind of what I mean.

Q253 **Katherine Fletcher:** So it relies on NHS IT.

Professor van der Schaar: This is again where AI can be very helpful, because as a clinician, I may not remember very well and I may not be able to integrate that information across time.

Q254 **Katherine Fletcher:** Understood. Professor Fernandez-Reyes?

Professor Fernandez-Reyes: As I said, we do not choose what the patient comes in with. The good thing about human training over 10 years is that I do not choose who comes to me; I have to be prepared for whatever comes. Of course, there are different things that come to you, because the populations are different. But what you describe is medicine. Medicine knows this. Over the hundreds of years that medicine has evolved, that is what the medic does. They look at people's risks, and all clinicians over time got very good at saying—

Q255 **Katherine Fletcher:** They are developing their own AI model.

Professor Fernandez-Reyes: Yes, but what is very difficult is to translate that into an expert system, or AI system. To summarise, the current context of AI is that if you have an overworked GP who has to assist people in five minutes, you can provide a reasonably good enough summary of where that person is. They are going to come in, and I need to look in five minutes, maybe do a test, and change their life. The AI has translated that. That clinician is overloaded.

Katherine Fletcher: That is the opportunity. Very finally, the bed and pandemic modelling must have been very bleeding edge, and was very important at the time, so that you for that work.

Q256 **Dawn Butler:** These are fascinating discussions. I was very much on board with everything that you were saying about AI in preventive medicine, precision medicine and personalised care. I think we are moving more towards personalised care, and AI could help with that. During your discussion, I thought about the risks of insurance companies, for instance, identifying a person using this technology and saying, "This person is more likely to develop diabetes, so their insurance premium is going to go up." How do we prevent that abuse?

Professor van der Schaar: That is a fantastic question. I lived in the United States for 18 years, and my mother and my grandmother died young of breast cancer. I was terrified to do a screening test for BRCA, because I was afraid that if I ended up having it, I would not get treatment. Luckily, I do not have BRCA, but the reality is that that is where regulation is very important. We do not want people not to do tests that may be relevant and may identify disease early, both because of the cost to the society and because of the cost to the individual. That is very important. I feel very pleased and fortunate to be in the UK system, which would not allow this differentiation. That is very much the case in some cases in the United States. Fortunately, with Obamacare, that is not really a problem any more, even in the US. We do not want to prevent people

from having insurance and high-quality healthcare. We want to encourage them to discover diseases and high risks early, and help them in personalising their care, given those risks. We need to realise that not all risk is genetic, and sometimes certain risks are modifiable.

Q257 **Dawn Butler:** That is going to take some stringent regulations and frameworks to ensure that there is not that abuse. Ethics will come into this strongly.

Professor van der Schaar: I agree completely with that. It is very important as a society to realise that we should not prevent certain people—it is not their choice to have a particular mutation or be a particular risk. We need to be able to serve those populations as well. But here regulation is very important.

Q258 **Dawn Butler:** What does Palantir do? Do you know much about the organisation Palantir? What does it do? What do you feel about the role that it is currently playing in our NHS system? Can you talk about that at all?

Professor van der Schaar: I am quite concerned about that. I would like to see much more focus not only on certain applications and certain niche applications, but much more on this area being open to many more companies and on better regulation. What I saw early on in the pandemic was too much focused on some areas rather than other areas of AI. Again, I think that we need to think about utilitarian AI that can have the individual and the NHS system at the forefront. I really would have liked to see more diversity of companies in this space, rather than only a few players, and much more transparency about why certain commissions are given to certain companies, to make sure that the right technology is in place across the board. At times, I am underwhelmed, as a machine learning researcher, by what these companies produce in terms of AI. It is very limited and not very cutting-edge.

Dawn Butler: Thank you. Professor Fernandez-Reyes?

Professor Fernandez-Reyes: I agree with Mihaela. That regulation and open competition is good for the industry. But it is like a black box in a black box. Our AI systems are black boxes in their own right, and we have a problem with regulating that. Now, if you put it in that setting, as you say, it is like you are putting a black box on top of that black box, making it much more difficult.

On what you were saying about insurance, it is already happening, I am afraid, with the use of these devices. There are insurers that provide cheaper premiums if you increase your activity. It is interesting—it is kind of an incentivisation approach—but it can have ethical repercussions as well. Suppose a person has a problem with a disease and cannot be active; their premiums are going to be more expensive. There needs to be interest in regulation of the insurance system, but that is a big task.

Dawn Butler: It is huge, yes.

Professor van der Schaar: I want to say one more thing. I am an academic, but I am also a big believer in open-source software, and I think it would be wonderful if, rather than having one company—Palantir or another company—the Government really embraced this type of open-source platform where industry, as well as academics, are providing this type of technology, it can be inspected by anyone, there is an equal barrier to entry for people, and as a community we can build tools that are generally useful. Too little is done on open source. My own lab, which is relatively small, tries to put as much as it can in open-source software that is sufficiently reliable, and we really hope that people take these technologies and build useful platforms for the NHS.

Q259 **Dawn Butler:** I totally agree. This is my last question. You mentioned something that is dear to my heart: AI reducing profits and centring on people. Do you worry that large organisations and companies that make a lot of profit from medicines and tablets will try to find a way to sabotage this kind of thinking?

Professor van der Schaar: Yes, as a person, I worry a lot. That is why I kept saying we shouldn't focus only on drugs and drug discovery and we shouldn't focus only on the industry, but we should focus on the individual, as well as the NHS. We want to make sure that we improve the healthcare of the population, give the minimum amount of interventions and give them when they are needed. I am afraid that at times the data we collect—not only at one point, but over time as interventions are made—is biased. The more we intervene in a biased and not okay way, the more data will show that is the right way to go, because we will not have counterfactuals for what would have happened otherwise. So, yes, I do worry about that. A lot more research and careful thinking needs to take place there.

Q260 **Dawn Butler:** We took evidence from somebody who said, "Oh, if we think about the bias, we'll never get anything done."

Professor van der Schaar: I am afraid that if we don't think about the bias, we will get a lot of very bad things done, as the research has clearly shown. I think this is an underappreciated area and a lot more research and thinking needs to come into it. It is non-trivial. There is bias not only at the collection point, but in interventions, because, over time, certain people get certain types of interventions and others do not.

This is what I meant when I said "reality-centric AI". This may not sound cool, but when you play a game of Go, for example, it is very obvious, because the game is played with very clear rules. If I think about folding proteins, again, the rules are clear; they are physics and biology. But in this case, there are a lot of interlaced factors and incentives that are corrupting the data. When we use data where interventions have been made, we need to be very careful not to propagate wrong decisions.

Professor Fernandez-Reyes: It is true; I agree. The pandemic shows us—and we know—that we cannot do this without industry. We need to find the fair and proper channels, and the interdisciplinary stakeholders, to

regulate and through which the industry can make their contribution to the regulation. Manufacturing is a problem, okay? The pandemic has shown that, and we have a decrease in manufacturing. It is not only about AI. If I cannot really produce something, it will not be translated into benefits anywhere.

We need to find that balance, and it is important to take industry on board. It is difficult, but we should try. The bias is nothing new, because medicine has been generated since Hippocrates' time; these biases exist. Most evidence-based medicine is trying to deal with that in one way or another. It is impossible to do so perfectly, but it is about safety. At the end of the day, there is a bias—some minority groups will not be represented and therefore they will suffer in a safety event, for example. That is important.

Professor van der Schaar: Definitely I am also for industry. I am not saying that industry shouldn't be involved. I think industry is very aware of that, but I think we still need to be careful and make sure that, with industry, we are thinking carefully about these problems.

Q261 **Chair:** On the question of biases, which is clearly very important, you were talking about the potential for bias in human interventions—for example, clinicians having a fixed, preconceived idea as to who will benefit, and what kind of lifestyles and things might be prioritised. Is it possible that the more expansive use of AI may overcome some of those human biases?

Professor van der Schaar: That is a fantastic question. As a matter of fact, this is exactly one of the important areas of research that I do. I do a lot of work on exactly this particular area. I coined an unhelpful phrase, "quantitative epistemology", to allow for that.

I work a lot with clinicians and exclusively a lot of transplant doctors, both in the United States and in Europe and the UK. We know that, for instance, historically women have been transplanted much less. They were considered to be not strong enough—to be weaker. That is one of the many examples I have encountered, and it took decades for that bias to be removed.

Some variation in practice is desirable because we have different ages, genders and characteristics, but there is also what my clinical colleagues call undesired variation in practice, where there are biases and wrong judgments may be made. What we are developing, together with clinicians in three different continents, is exactly this type of technology that tries to see what is undesirable bias and tries to accelerate the adoption of technologies. In particular, we have written some recent work where we show if we had identified the technologies that we are building, which are actually quite complex machine learning models, probably ten years earlier—from the data—humans could have been transplanted for heart surgery earlier.

These are exciting technologies; you are right. But it is complex, because the data that we have encodes knowledge, expertise and positive biases, but also negative biases. Disentangling the two can only be done by building careful technology in collaboration with human experts and a diversity of human experts.

Chair: That's very clear.

Professor Fernandez-Reyes: But just because that comes from an academic environment, that doesn't mean that doesn't need to be regulated. That is important. Let's take a simple example. Let's suppose I am a GP in Chelsea, and the next day I have to be—I have never been there, but I have been called to be—a GP in Tower Hamlets. I have been working as a GP for 10 years in Chelsea, and then I go straight ahead to the Tower Hamlets population. As a medic, that is an issue. If these AI systems are done nationally or internationally, the concept of homogenisation in some aspects can allow that transition for this new clinician, by saying, "Hey, you were not paying attention to these things because you have a biased population set." It is augmentation of clinical care, so the individual benefits. It is clear that that will be a benefit in a lot of different settings, some minorities are represented in some settings and not in others.

Professor van der Schaar: It is not so much that we will build technology that will say, "Hey, you are biased" and will make the clinician aware so that they debug their own thinking process, but more or less that they will ponder different aspects. For example, just because we are women doesn't mean necessarily that we are weaker. It is about just letting them ponder and showing them related examples from cases locally and internationally. One of the challenges we saw was that in some centres they were seeing very few subpopulations with certain characteristics, so they didn't have as much experience and were afraid, for instance, to act. But the overall information throughout the globe existed. What I am after is definitely not replacing clinicians, but providing them with tools such that they can ponder different choices that they have, provide them with evidence and then let them, of course, make the final decision together with the patient.

Q262 **Dawn Butler:** I think you mentioned a black box on top of a black box when we were talking about Palantir. You mentioned how worried you were that we have put all our eggs in one basket with this one company. The Palantir black box could be full of all kinds of biases that we don't know about, right?

Professor van der Schaar: That is actually a problem. That is why we need a lot more open source. As a country, we need to invest in open source software such that many different companies and groups can inspect the data and the technologies that are built, and we can get stronger together.

Professor Fernandez-Reyes: Is it a necessary condition but not sufficient. As the MHRA has already hinted, we need the tools to deal with

that, because I can make the code open source, but code changes. There are a lot of technologies we need to implement this. The system is changeable; systems are sometimes monolithic, sometimes they are adaptive. I do not think open source will solve the problem completely. When I develop a drug, the MHRA, yes, sees everything, but that drug is not open source. I don't know the source components of the Pfizer vaccine, for example—the nanotechnology—but the regulator knows. That is more efficient than me making the components of that thing open source, because it might be a necessary condition, but it is not sufficient to achieve what you are saying. We have had open source for many years. Journals thrive on open source and so on but—

Professor van der Schaar: You misunderstand my point. It is not about the AI software that companies like Palantir and others are using internally, and getting that scrutinised. What I am saying is that I would like to see a lot more investment by the Government in open source AI software, which can be used as a platform for other companies to build their knowledge on, so that it is transparent to everyone and can allow many companies to flourish, rather than putting everything in a company that does not share their software.

Chair: Very good. We are very grateful for that discussion. The purpose of the inquiry is the future governance of AI, so these questions are very relevant. That concludes this meeting of the Committee.