

Science and Technology Committee

Oral evidence: [Algorithms in decision-making](#),
HC 351

Tuesday 16 January 2018

Ordered by the House of Commons to be published on 16 January 2018.

[Listen to the meeting](#)

Members present: Norman Lamb (Chair); Vicky Ford; Bill Grant; Darren Jones; Stephen Metcalfe; Carol Monaghan; Neil O'Brien; Graham Stringer; Martin Whitfield.

Questions 208 - 293

Witnesses

[I](#): Dr Dominic King, Senior Staff Research Scientist and Clinical Lead, DeepMind Health; Dr Ian Hudson, Chief Executive, Medicines and Healthcare products Regulatory Agency (MHRA); Professor Harry Hemingway, Farr Institute of Health Informatics Research; and Eleonora Harwich, Head of Digital and Tech Innovation, Reform.



Examination of witnesses

Witnesses: Dr Dominic King, Dr Ian Hudson, Professor Harry Hemingway and Eleonora Harwich.

Q208 **Chair:** Welcome, all of you, and thank you for being here. I would like you each to introduce yourselves in a moment.

With a panel of four, if everyone answers every question we will be here all day, so will you answer only if you want to add something to what others have said and try always to keep your answers succinct, so we can get through the agenda?

Let us start with introductions—we will start with you, Dominic.

Dr King: Thank you. My name is Dr Dominic King. I am the clinical lead at DeepMind. My role at DeepMind is to ensure that all our projects in health are led by clinicians and patients and backed by strong clinical evidence. Before joining DeepMind two years ago I worked as an academic general surgeon at Imperial College London. I was in clinical practice for nearly 15 years and led programmes of research in patient safety, particularly around the lack of good technology in healthcare, and in digital health.

Dr Hudson: I am Ian Hudson. I am chief executive of the MHRA. Our interest is that algorithms and software used for medical purposes fall within the definitions of medical devices. We regulate medicines and devices, and when algorithms, software and so on are used for medical purposes they need to be CE marked, either through self-certification or through notified bodies. We oversee the notified bodies, and we are also responsible for monitoring safety and taking action if there are safety issues in the marketplace.

Professor Hemingway: My name is Professor Harry Hemingway. I am from University College London. I represent no specific organisation, but I would draw the Committee's attention to Health Data Research UK, which is the next iteration of the Farr Institute. The UK is special internationally in working across universities, the NHS and industries in driving forward the benefits from algorithms that can be derived from the very special NHS data. Research is my focus.

Eleonora Harwich: I am Eleonora Harwich. I am the head of digital and tech innovation at the Reform think-tank, which focuses on ways of delivering value for money within public services. I have recently published and co-authored a paper on the applications of artificial intelligence within the NHS.

Q209 **Chair:** Perhaps we could start with those of you who want to answer this giving a snapshot of how algorithms are being used in healthcare now and in the future. Where are the biggest potential gains to be made?

Eleanora, you have just produced a report, so perhaps we could start with you, and others can chip in if they want to.



Eleonora Harwich: We were chatting in the corridor about the distinction that needs to be made between the newer types of AI and the older types: algorithms that might not be classified as artificially intelligent or the older expert system type of algorithms that have been more widely used within the NHS.

As for what they can do and their potential, one feature is around better prediction and detection. There is also that of considering how individuals can better self-care and potentially live more healthily through the use of algorithms.

There is also the automation of different back-office functions—for example, better scheduling of operations. On detection, different things can be done—for example, in image recognition, which is an interesting field that is currently developing in machine learning. One example is improved diagnostics when looking at mammography scans and better detection of breast cancer.

There is also the use of machine learning for unstructured data, looking at speech recognition devices. I recently read an article on how that is being used in Denmark.¹ When people call and someone is having a heart attack, they have an intelligent virtual assistant that parses language and dispatches units when it detects that someone is having a heart attack. That has an increased accuracy rate compared with how humans can classify those calls and then dispatch.

Dr King: Doctors and nurses have been using algorithms in the NHS for many years, as Eleonora says, for identifying the deteriorating patient in a hospital, for example, or for getting guidance on a heart trace, ECG trace or scan. It is widely recognised that many of these lead to many false positives and false negatives, which lead to problems with underdiagnosis, overdiagnosis and misdiagnosis. Although there is real excitement, hence why we are sitting here today talking about the role of artificial intelligence and improving these algorithms, I, as someone who has been working in the field for a number of years, would struggle to identify one example of the kind of advanced AI that we are building at DeepMind being deployed at scale across the health system and having robust clinical evidence demonstrating its positive impact.

It is really great that we are having this conversation now, before wide-scale deployment, because lots of issues need to be ironed out, but at the moment there is quite a lot of hype. The onus is on us working in this space to translate the interesting research that is going on at the bench into meaningful impact for patients and at the bedside.

Q210 **Chair:** Is there the potential to achieve the prize of better outcomes and reduced cost?

¹ Note by witness: The following paper addresses the cost of not sharing data in healthcare, *The other side of the coin: Harm due to the non-use of health-related data* (International Journal of Medical Informatics 97 (2017) 43–51)



Dr King: Absolutely. You only have to open a newspaper to see the struggles of the NHS, but this is a common tale across different health systems. There are already early promising signs, and we are very excited about the potential, as Eleonora says, for using these algorithms to detect patients in earlier stages of deterioration, to identify disease at an earlier stage in scanning, and to identify new drugs through computational methods.

These technologies have an absolute ability to transform healthcare, but we need to remember that if you walk into an NHS hospital today the likelihood is that your prescriptions will be done on paper, the doctors will be communicating with pagers, and referrals will still be coming in by fax machine. These things need to be sorted first before—

Q211 **Chair:** I will come on to what the NHS needs to do. Do you have any final comment about the potential for algorithms in healthcare?

Dr Hudson: To echo my colleague's comments, I think there is huge potential to transform the health service, building on where we are and the use of algorithms now—whether it is a defibrillator, risk scores or whatever it is.

It is not without risks. It has to be a steady, stepwise journey, thinking through the issues as we go through. Plenty of issues will surprise us along the way, but we are on a journey with real potential.

Q212 **Chair:** The next question is: what does the NHS need to do to realise the potential? What is the challenge that the NHS faces?

Professor Hemingway: The first challenge is to understand the deep benefits of the data that the NHS has. It is often stated—it has been stated for decades—that, with the advent of artificial intelligence and associated technologies, leveraging the potential for that rests on understanding that benefit. By those benefits, I mean the fact that there is cradle-to-grave information—that nearly everybody in the country is registered with a GP who uses a structured electronic health record. That is crucial.

On the potential benefits, there is not an aspect of what we do that improves health that is not potentially informed by or improved by artificial intelligence. I underscore the point that this Committee inquiry is timely. We are at an inflection point, where the number of publications in the medical sphere using deep learning and other approaches is flicking upwards, but they are all at an early stage. They are at an early stage in asking, "Does the computer do better than the human?" That is extremely important, but it does not answer the questions, "Did the human patient benefit? Did they live longer? Did they suffer less? Were costs saved in the NHS?" This constant theme of evaluation is vital.

Q213 **Chair:** Following on from your earlier comments, is it the case that we have a particularly special asset in the NHS, because of the data and the whole system? Does it put us into a very good position to exploit, in the



best possible sense of the word, the potential for this technology?

Professor Hemingway: I would quote from the Life Sciences Industrial Strategy, from Sir John Bell, which uses the phrase “globally dominant.” It says that it is the combination of having—for example, we have heard about imaging or pathology at scale—with the patient context, the longitudinal, over-time assessment of “Who is this patient? How did they get there? What are their outcomes?” That combination allows us to say that we have a duty to return the value of these technologies to our patients and citizens. I think there are economic benefits in so doing.

Q214 **Chair:** Dominic, you described the way the NHS too often operates with faxes flying around every day of the week and so on, with the sense that it might be caught in a bit of a time warp.

Dr King: Yes.

Q215 **Chair:** What does the NHS have to do to prepare itself for taking advantage of the technology?

Dr King: We are in an incredibly good position. As a consequence of our world-class universities such as University College London and Imperial and Oxford, and companies such as DeepMind, we arguably have the greatest concentration of artificial intelligence talent and expertise anywhere in the world. As Harry says, we have health data that can be used to advance artificial intelligence.

There are some other critically important success factors. World-class clinical expertise is required to think carefully about how these technologies are implemented safely and effectively. Considerate regulators are thinking through the important issues that must be resolved, in a way that I am not sure many other countries are.

Q216 **Chair:** You are conscious of who is sitting next to you.

Dr King: Yes—of course.

We are having a very mature conversation publicly, if we consider the Wellcome Trust and understanding patient data and the various efforts that are being made—and the fact that we are sitting here. As I was saying to Harry before we came in, when I finished medical school back in the early 2000s, it was right at the height of the start of laparoscopic surgery. Everyone was starting to do it. We were not having conversations like this before the technology was introduced, and there were lots of widely documented problems that were a consequence of that. This kind of careful conversation that we are having now is very important.

Q217 **Chair:** But is there a need for a significant investment in the digitalising of the NHS to get from where we are now?

Dr King: Yes. We have underinvested in technology in healthcare. It surprises me that, 15 years ago, when I started as a junior doctor, I was



HOUSE OF COMMONS

given a pager, and that is still the case in pretty much every NHS hospital.

We are making progress on the slow journey to becoming paperless, and that will only be of benefit to the clinicians who work in hospitals and who currently spend well over half their time on administrative tasks and non-direct care, as well as to patients, who get frustrated by not having their data shared effectively and by not being able to book appointments online like they do for a restaurant or hotel.

Q218 Chair: Dominic, the Royal Free agreement with Google DeepMind gave free access to 1.6 million personal identifiable records without patient consent. Can it ever be right for a deal of that sort to be entered into without consent having been secured?

Dr King: It would be useful to give a little background to that partnership. About three years ago, a team of clinicians from the Royal Free approached DeepMind, interested in the problem of acute kidney injury, which is sudden damage to the kidneys. The problem kills about 40,000 people a year just in England, and a quarter of those cases are entirely preventable. We entered into a partnership with the Royal Free to develop a modern digital technology that actually does not involve any—

Q219 Chair: I understand the value of the partnership. What I would like you to deal with is whether you recognise that there were failures in the deal that was done and whether that needs to be addressed for the future.

Dr King: Absolutely. The findings of a very lengthy investigation by the Information Commissioner's Office were directed at the Royal Free, but we are partners and take our share of blame for those failings. The two issues that were raised particularly concerned the use of patient data for the safe testing of software, and then the failure of us, in the Royal Free, to provide adequate information to the patients we were looking to serve about how their data were being used and processed. I absolutely think that patients have the right to know how their data are being used and processed.

It is important to say that this was not an AI research project; it was for direct care. On your question on the issue of consent, the reality is that, for direct care, most NHS organisations do not ask patients to opt in or opt out of direct care services such as the software that runs CT scanners for lab systems, because it would be unsafe to do that. That decision was taken by the Royal Free, as the controller of the data, around the issue of consent.

Q220 Chair: As a company, do you recognise the importance of the NHS benefiting financially from deals with the private sector that might realise profit for the private sector? Given this prized asset of data that the NHS holds, the NHS and, indeed, patients must be seen clearly to benefit from those arrangements.



Dr King: Absolutely. As I said, there are a number of success factors, and the prize on offer is very substantial—I mean that in terms of clinical impact. If revenue is being generated, in all our partnerships to date we have entered into very useful conversations with our partners that will ensure that all the partners involved realise value from that.

Value can be generated or delivered back in many ways, from enhanced capability building to free product and payment for data—I am not suggesting that I think that is a good idea; it is a model that works in other countries, not necessarily in the UK—or some share of revenue or IP. From DeepMind’s perspective, we are happy to talk about all those considerations, both with our partners at a local level and on an NHS level.

Q221 **Darren Jones:** I want to pursue the questioning around the value of data. First, how would you describe the value difference between the algorithm and the data that train and drive the algorithm? How much is an algorithm worth, in value terms, without the data?

Professor Hemingway: Let us start with what we know. Before AI, we had thousands of algorithms published in the medical literature. They are easy to develop: things that would predict our risk of cardiovascular disease and other diseases. Those are easy to develop.

Imagine a pyramid, with thousands at the bottom. The next step up is, “Do they work in practice?” Are they valid? Very few of them are actually valid. That is getting into your value question. The next step up is, “Are they used in clinical practice?” A tiny proportion of those thousands are used in clinical practice. At the apex of the pyramid is, “Are they useful?” Do they return value to individual patients in improved patient outcomes or better population health?

We can learn from that extensive history and we can say that, with AI, we have an opportunity to do something different—to put the development, validation, use and impact of these algorithms into the context of a robust, evaluative framework. I think that is a real opportunity, and we would bring together multiple stakeholders to do this. Health Data Research UK is a national forum to do that.

In answer to your question, I think they are inextricably linked, and the NHS should not lose sight of that.

Dr Hudson: I agree. You cannot separate the two. You need data to be able to develop the algorithm, and you need data to run the algorithm on. If you just have data or if you just have the algorithm, it will not work. The two must go together.

Dr King: Algorithms are valueless without data.

Vicky Ford: And data are valueless without algorithms.

Dr Hudson: The two need to go together.



Dr King: The science of algorithms has a powerful impact on their eventual effectiveness. In healthcare, the difference between being 96% accurate and being 98% or 99% accurate is huge. You cannot underestimate the science that is going on at our universities and, for example, at DeepMind. We have a team of 700 people, including some of the world's leading engineers and research scientists, and they are all motivated about getting to that 99% accuracy, which will, as Harry says, make the algorithms useful and clinically impactful.

Data are the lifeblood of these algorithms. Both are incredibly important.

Q222 **Darren Jones:** Do you think it was right for the Royal Free London NHS Foundation Trust to offer its data to you at DeepMind for free?

Dr King: It is very important to get across the difference between what we are doing at the Royal Free, which is a direct care service, not a research project. It does not involve AI. We feel strongly that improving the digital maturity of hospitals is an essential first step to the wide-scale application of artificial intelligence. Our partnership at the Royal Free was not about us taking data, doing research projects on it and developing algorithms to sell in other jurisdictions. It was about delivering a direct care service, which hundreds of other software companies already provide to the NHS.

In our other research partnerships—for example, with Moorfields Eye Hospital, where we are looking at eye scans; at UCLH, where we are looking at radiotherapy planning; or at Imperial, where we are looking at mammography—we have discussed the opportunities around bringing value back to our partners. Those are very different projects.

Q223 **Darren Jones:** When we say “bringing value back to our partners,” what is that conversation with the NHS partners like? What does that mean?

Dr King: We have spent a lot of time thinking about what value looks like. It is important to say that these projects are at a very early stage. Although there are promising early signs, we need to prove not only that the algorithms work but that they are clinically impactful. That will probably take a couple of years of clinical studies, as you would imagine for any pharmaceutical or medical device.

If that is ultimately successful, the terms of our agreement are that, in terms of the value we give back to the partner, we have invested a substantial amount of resource in enhancing the datasets that we are processing to ensure that they are ready for artificial intelligence, so that they can be shared with other research partners. For example, we would give any eventual product back free to Moorfields Eye Hospital so that it can use that algorithm, either for a specific set number of years or in perpetuity. If we ever deliver any revenue, we think about how we would share that back with the partner.

Q224 **Darren Jones:** Do partners share any of the intellectual property rights with DeepMind?



Dr King: Currently with our existing partners, we have talked more about revenue share. One thing that it is important to recognise with the algorithms that DeepMind is building is that they are generalisable algorithms. We are using the same algorithms in many projects outside healthcare as we use in healthcare. This is something we call artificial general intelligence, as opposed to narrow artificial intelligence. That has implications, from our perspective, about the sharing of intellectual property. That is not to say that we do not see that there should be some type of share in any financial benefit that is gained from these types of projects.

Q225 **Darren Jones:** To summarise so that I am clear, you are saying that you have general AI or algorithms, which you can use and train using NHS data, but you get the value from that by using non-NHS-use cases.

Dr King: No. We are taking the progress that we have made in many other areas. Algorithms are composed of lines of code that have been effectively developed in many other areas. Effectively, they are general learning algorithms, which can then be deployed in many different applied areas of operation, one of which is health.

Q226 **Darren Jones:** Do you see this value conversation changing in the future beyond, “We’ll do stuff for you because you can’t do it yourself, and we might share some revenue”?

Dr King: Yes. It is an important time for us to have this conversation. As everyone on the panel and everyone here recognises, there is huge value in the data, and it is important that we have a real conversation. Many of the partnerships that we are organising are at a local level and are discussed directly with Moorfields, with some discussions more centrally, but we are very much open to whether a more central conversation needs to happen. We recognise Sir John Bell’s recommendations. It is the right time to have these conversations.

Q227 **Darren Jones:** What about others on the panel? I have a commercial background, and in my life before being an MP it would have been very odd for it to have been in my client’s interest or, indeed, in their commercial interest to offer such a valuable source, technically, for free, for value definitions that in my view are not very clear. Do other people on the panel agree? Do we need to think about value differently in the future?

Dr Hudson: The only comment I would make is on the public health value of data. We have a particular interest in the MHRA in that we house the clinical practice research datalink, which is a set of anonymised primary care records that can be linked with some secondary care data. It has been used for many years. It is a very large dataset, and it has been used for observational studies and pharmacoepidemiological studies, addressing some pretty significant public health challenges—studies such as refuting the link between MMR and autism, studying the safety of pertussis vaccine in pregnancy and looking at statins in the



HOUSE OF COMMONS

marketplace. There is real value, as in public health value, in addressing some very important public health issues. We should remember that value in the discussions. I want to talk not about the commercial side—that is outside my remit—but the public health value.

Professor Hemingway: It is an early stage of this conversation. On value, we are talking about commercialisation of something where we are not sure what it is or whether it works, so this is somewhat premature. It is an appropriate discussion, but it is somewhat premature.

The debate about how you capture that value for the NHS and retain it in the NHS is extremely important. In addition to financial value and patient outcome value, the NHS has other values in the mix, which are equalities. We know from other sectors that, if one does not monitor algorithms, they may bring in things or perturb things that we care about regarding the social agenda or ethnic equalities. The NHS needs to consider how its purpose is advanced through this.

Q228 **Chair:** It should not just be left to individual trusts necessarily. Do you think there should be some sort of framework?

Professor Hemingway: I strongly think that there should be some kind of framework, and the UK has an opportunity to do something stand-out, to which others in the world might look. Others in the world look at the way we do things in NICE, in MHRA and so on, and we have an opportunity to do likewise here.

Q229 **Neil O'Brien:** Could the NHS or, for that matter, other Government Departments be doing more with the data they hold to boost innovation, the growth of the AI industry and the use of technology? This is almost the same question, but from the opposite angle. What more could we be doing?

Sometimes people have said to me that it is quite difficult to sell into the NHS, even for things such as data science. What more could we be doing in the NHS or in the rest of the Government to promote the use of our data and the joining up of data with algorithms and AI? I would be particularly interested in the views of Eleonora and Dominic.

Eleonora Harwich: One of the important things to recognise is that the NHS has quite a lot of trials with companies that develop algorithms or that use AI. There are several points that I wish to make from the previous questions, which relate to what you are saying.

There is a point to be made on the value of data. Data have value, but what really brings value in patient outcomes and ensuring that algorithms do not end up having unintended consequences of further entrenching healthcare inequalities is having high-quality data. We should be thinking about that, in that it is not only about the expansion or scaling out of these tools; it is also about ensuring that the current basis of data reflects the best policy.



HOUSE OF COMMONS

Q230 **Chair:** There is a job to be done there in improving quality of data across the NHS.

Eleonora Harwich: I think so—definitely. Things are being done. There is the standardisation of clinical codes, which are going to be replaced by a standard system. Before, that was not the case. That meant that, depending on the IT provider that trusts or GPs used to have, they had different categorisations of disease, which meant that it was difficult to link up data.

Q231 **Chair:** That presumably makes all the historical data of limited value, to some degree, if they have all been coded differently.

Eleonora Harwich: I mean, IT—gosh—

Chair: Harry—quickly.

Professor Hemingway: The historical data are a treasure trove.

Q232 **Chair:** But if they are all coded differently by different providers, is that—

Professor Hemingway: There are ways around that. The question I am most commonly asked is, “Are the data fit for purpose?” Time and again, we see that they are better than we think they are.

Q233 **Neil O'Brien:** So, we are in the process of standardising the data. What other things might be helpful in catalysing the use of big datasets? Is there a good portal for people who want to do these things to interact with Government? Is there some other thing that we should be doing but are not doing at the moment?

Eleonora Harwich: I will make one last point. One of the things that are not currently being done—I might be wrong and you might correct me—but that would have huge value is looking at potential new sources of information, for example from trackers, sensors, wearables and so on, which would provide a different type of data that the NHS currently does not have: a continuous understanding of heart rates or other vital signs. It would also give an interesting perspective of our understanding of people’s health. Some of this would give us information about people when they are healthy, not just when they present.

Q234 **Neil O'Brien:** That would generate more data.

Eleonora Harwich: Yes.

Professor Hemingway: Health Data Research UK, which is launching this year—director Andrew Morris; chair Graham Spittle—is led by a funding consortium of the Medical Research Council and nine other funders. That will, across multiple sites in the UK, which will be announced in three weeks’ time, drive forward four research programmes, one of which is actionable analytics, including AI. That is a concrete, nationally and internationally visible place in which the NHS, academia and the wide range of industries—plural—can come together.



HOUSE OF COMMONS

To emphasise that last point, it is not just very clever companies such as Google, DeepMind and others in tech; it is biotech, it is UK and international pharma, and it is small and medium-sized enterprises.

Dr King: Over the past few decades and longer in the UK these collaborations where we think about these issues proactively have led to the development of transformative technologies, from CT scanners to MRI scanners to many of the pharmaceuticals that we and our families take. There is a long track record of us getting this right.

On the datasets, we all recognise the importance of the underlying data, and I absolutely take Harry's point that the data do exist. From our perspective, there is a lot of work to be done to get them into a machine-readable, AI-ready format for research. Eleonora talked about issues around interoperability and common standards. We have found with our work at Moorfields Eye Hospital, for example, that, once the partnership is signed and data are transferred from the controller to the processor, our research does not start the next day. It took many months of cleaning and labelling the data.

On our algorithms, just to give some background, there are two types of artificial intelligence: supervised learning and unsupervised learning. We are using supervised learning approaches at the moment. That effectively requires some type of training material. The NHS does not have datasets that have been meticulously labelled, pixel by pixel, for training material for artificial intelligence companies. That is what we have created in partnership with Moorfields Eye Hospital. It now has, unarguably, the leading retinal scan dataset in the world, over which we have taken no proprietary ownership. We have said, "We want you to be able to share this with others."

Q235 **Neil O'Brien:** What would you like to see happen that is not happening at the moment?

Dr King: Better education and investment in what it takes to get these datasets ready, so that they can be made available to a wide group of people. We recognise that we have a very privileged position where we can undertake the investment to get that dataset ready, and we would never like to be in a position of saying, "We are the only people who can look at these retinal scans." Within those retinal scans are the early signs of people going blind.

As far as we are concerned, any organisation that can improve on our algorithms should have absolutely every right to do so. I think the NHS should work towards developing similar high-quality representative datasets, and by "representative" I mean that it is not good enough to have one eye hospital or one hospital's dataset. What happens at St Mary's in Paddington, where I used to work, is very different from what happens in the north-west of England in terms of demographics. We need large, representative datasets that are appropriately de-identified,



HOUSE OF COMMONS

cleaned and labelled, which are then ready, with the appropriate consent and approvals—

Q236 **Chair:** That is a pretty massive undertaking across the system.

Dr King: I think so.

Q237 **Stephen Metcalfe:** Once you have cleaned the data, you presumably add some data about data, or metadata—

Dr King: Absolutely.

Q238 **Stephen Metcalfe:** Are you doing that to an established, recognised standard that is now adopted across the whole UK, perhaps Europe and beyond? Otherwise, is there not a danger that different people will be using different datasets, and you will have to do it all again at some other point in the future?

Dr King: Yes. We have had a big problem with a number of common data standards in the past, both for direct care and for research, where data have not been standardised. The whole community, from academia to the commercial sector, now recognises that it makes no sense for data not to be in a structured format. A lot of progress has been made in that space. We will absolutely work to the latest established frameworks and common standards.

Q239 **Stephen Metcalfe:** But who establishes that, and who is in charge of coming up with those standards?

Dr King: Different groups of people. For example, for much of the direct care work we are doing, there is a common standard called FHIR—I sometimes forget what that stands for, but it is fast healthcare interoperability resources. That grew out of Boston Children’s Hospital, where a team of academics and IT specialists got frustrated at the lack of common standards and effectively built their own. It is now widely accepted across the sector as the common standard.

Q240 **Chair:** But this needs national and indeed international co-ordination.

Professor Hemingway: Absolutely. There is something called the international classification of diseases. How it is applied consistently in the NHS and across countries requires scrutiny, let us say.

Chair: Eleonora, I think you wanted to come in here.

Eleonora Harwich: No—I was just nodding away—sorry.

Professor Hemingway: Could I pick up on Neil’s question about what the NHS could do differently and what it could do that is new, and to pick up on Darren’s point about not being able to do it yourself? I think that the NHS, with ambition and thinking about its future and its sustainability in particular, needs to say, “Under what circumstances, can a public sector body deliver patient benefits in this?” As I say, from soup to nuts, in anything that we can do to improve health, AI has a potential role. Can



the NHS do it now? Well, guess what—the staff were not employed for those reasons. Should it be seeking public sector and appropriate industry collaborations to grow its own AI interventions? I think there is an important question there.

Q241 Vicky Ford: On this issue of standards and trying to have interoperable datasets in different hospitals and different countries, which is obviously an industry-led process, is there a regulator role in developing those standards? Does the regulator feel involved in the ethics being considered from a regulatory point of view?

Dr Hudson: Our role is the regulation of the algorithm, if you like, rather than the data that it is run on. There are others that are involved a bit more in the data governance side of things. In terms of the standards that we would expect, this is an area where there are some international standards. The International Medical Device Regulators Forum has produced standards in developing algorithms and what you need to see. We are in the process of developing a guidance document with NHS Digital, which should be going out in the not-too-distant future, on the high-level requirements. Quite a lot of work is going on and has already happened in defining standards and in what you need to do to get your algorithm CE marked and to get the data that are needed.

Q242 Martin Whitfield: How good is the NHS at sharing data among itself? Do you know? Can you comment on that—politely?

Dr King: If we divide this into direct care services and research, for direct care services everyone in this room will have been frustrated at some point. The reality of modern healthcare is that you move between primary and secondary care and between different hospitals.

I experienced this a couple of years ago. We had a premature son born, we went to four hospitals in London and we received excellent care from each, but there was absolutely no joining up of the records in any way whatsoever. There is not a canonical record of his care.

That is a big problem for delivering safe, high-quality healthcare. Everyone recognises that as an issue, and a lot of investment and effort are going into trying to improve that, with FHIR standards. The move to paperless will undeniably help.

Similarly, for research, as to the idea that we could go to a central body, whether in my past academic life or in my new life, and ask for every MRI scan of a right knee from the past two years from the NHS, those datasets do not exist—although there are some areas where they do.

Bringing together data from different geographical areas with different ethnic and diverse make-ups is really important in building those representative datasets, which then prevents some of the issues that I am sure we will come on to discuss, on bias in algorithms. Data sharing has a long way to go to get to a point where we would all be happy with it, either as a patient or as someone interested in healthcare research.



Q243 Martin Whitfield: I presume that everyone is in reasonable agreement with that synopsis. If we take the direct care element and we consider informed consent, when you attend the hospital, you have surgery proposed to you. Although the concept of informed consent is perhaps quite difficult, it is always attempted, so that the patient understands what is going to happen and the potential risks. When we move to that patient's data, the concept of consent changes fundamentally. Do you have any comment on the giving of consent to the data, before I explore it a bit more? What is your view on a patient's consent to their data being used and on the quality of the consent that is given at the moment?

Professor Hemingway: Is this under direct care at the moment?

Q244 Martin Whitfield: No—I am using informed consent for specific tasks. I am looking at the data that are derived from direct health and the data that sit within those datasets. What is your view of the quality of consent that is given at the moment—or that is even discussed with the patient?

Professor Hemingway: We understand from the understanding patient data initiative led by the Wellcome Trust that citizens of this country are unclear about how their data are used, across trusts in direct care and for research under different headings.

As you know, the National Data Guardian's report is proposing, after consultation, the opt-out. Consent and the nature of consent are at the heart of what we are discussing today. To see that in the context of the benefits or potential benefits that come from data, algorithms and the use of data, through to the value and the ownership of the value in that, is a wider societal debate.

Clearly, the traditional models of consent that have grown up over decades are challenged by the current situation. Let me give you one example from the setting of whole-genome information by Genomics England, which is an NHS transformation project. The understanding of the value of my personal genome—my data—is intrinsically related to the understanding of that information from other people. Distinguishing what is direct care and what is research is ever more blurred, notwithstanding the issues of dynamic consent and understanding how the uses of the data may reasonably and for patient benefit change over time.

Dr Hudson: Clearly, there must be appropriate governance and oversight in place. It also depends on whether the data are anonymised data or identifiable data. Those are quite different circumstances. What are they going to be used for? I do not think there is a simple answer to that.

Q245 Martin Whitfield: Which there would not be. Perhaps we can explore that.

You suggest that there is a difference when the data are anonymised. Would you like to explain what the value is and what the depth of consent that is needed from an individual should be where their data are



going to be anonymised? Is it the case that, once they are anonymised, the consent should come from the holding partner or the partners in it? Do you have views on that?

Dr King: Only to say that the UK, again, has been a thought leader in this area. You will all be familiar with Dame Fiona Caldicott's recommendations and work. Broadly, it is seen that anonymised or depersonalised data can be used without individual-level consent for research. That strikes most people as being sensible. We are working on a project at DeepMind called "verifiable data audit," which aims to provide partners, regulators and, ultimately, patients with cryptographic proof of how their data are being used and processed. Our security engineers do not like to call it "blockchain," but that type of technology allows people to see how their data are being used. Dynamic consent processes, as Harry suggested, is where we would all like to get to, but it is very difficult at the moment.

Professor Hemingway: To be clear, the ability to analyse anonymised data has demonstrably led to multiple health benefits. Preserving that is—

Q246 **Martin Whitfield:** The data are not just financially but intellectually invaluable, and the value of the data exists because the amount of it is so large and they are not any one individual's data—it is the cumulative effect of the data. The data will also be used for things that were never under consideration when the database was created. Is it a case of where and on whose responsibility this rests? Is it a case of educating the individual patient, society and indeed all users of the NHS in the value of a pretty blank consent that they are giving, simply because, as you said earlier, the public health value of that accumulated data is so great?

Professor Hemingway: Yes, it is about education and value, but why do people share intimate details in public on social media? It is because they get some immediate return. Let me give you an example of immediate return. If any one of us walks into a healthcare organisation with some problem, we might legitimately ask, "Tell me about me. Tell me about other people like me. What are we likely to suffer? What drugs and procedures are we likely to get?"—in other words, any questions about me. I want to know that the system, the NHS, has learned from other people like me. Right now, it is not possible to answer those questions. If we could answer them, I think we would return interest, benefit, understanding and engagement with patients.

Q247 **Martin Whitfield:** The greater value of the data and the pooling of the data are enough to forgive specific consent to the nth degree, because the whole—

Professor Hemingway: It is a proportionate argument. We absolutely need to protect privacy. The real harms that can be done could potentially be through re-identification itself, on the one hand. What we do not see, and what does not make the headlines—I am not going to



HOUSE OF COMMONS

attribute this quote—is the plane-load of individuals who are falling out of the sky every day: people who are dying because we are not sharing the data.

Chair: Martin, are you nearly—

Q248 **Martin Whitfield:** I was just going to come on to the last bit, to Dominic, on that. I know that we have discussed the data protection events but, regarding consent, what has been learned about your experience?

Dr King: Mixing up direct care and research a bit, the key finding from the ICO was about informing patients and the public adequately about how their data are being processed. As other panellists have already said, we know from the Wellcome work that patients and the public generally have a very poor understanding of how their data are currently processed.

What we do at the Royal Free and our partnerships with Moorfields, UCLH and Imperial College London had a transformative impact on how we try to tell patients about how their data are being used, from leaflets in the wards to standing on streets in Somerset or in shopping centres telling people about how we are going to be working with hospital partners to process their data and about updates on websites. That is a really important thing that everyone in this space should be doing more.

Q249 **Martin Whitfield:** Do you think you have moved to—using the phrase that is often used—a gold standard for that information on how their data are to be used?

Dr King: I would not say a gold standard—there is always learning to be done—but we spend a huge amount of time engaging and working with patients, patient groups and clinical groups. We are about to go live with streams for the product that we have at the Royal Free and a number of other NHS hospitals over the next few months. There are very substantial differences in the process that we are going through in the next couple of months compared with what we went through a year ago. Lessons were learned.

Chair: Everybody is going to have to be disciplined, as we are very tight on time. Eleonora can come in quickly, then Carol, and then Darren.

Eleonora Harwich: I will quickly echo one point about the cost of not sharing and the importance of education. It is important not only to educate the public but to have a conversation—it might sound a bit clichéd to say it—about a renewed social contract on how we understand data. That concerns not only the cost of not sharing but the potential cost to others of opting out and how that can skew samples and our understanding of population health.

My second point is to echo what Dominic said about the interesting solutions that can be found in new types of technologies such as



distributed ledger technologies. Their real benefit—they are not going to be a panacea for issues of interoperability and that type of thing—is about detecting where data tampering can happen and creating a list of decentralised access controls, so that people understand how data were used and for what purpose. It is also interesting how you can codify data protection legislation within a smart contract on distributed ledger technologies.

Q250 Carol Monaghan: I will try to be quick, Chair.

Dominic, you started by talking about your experience with your son. To me, we are talking about two different types of data: identifiable and anonymised. I understand—you can correct me if I am wrong—that NHS Scotland is far better at sharing identifiable data among different healthcare providers. Are there lessons that NHS England could learn from NHS Scotland, or does the vast variety of trusts make that impossible?

Dr King: I can say a little bit about this, having spent some time with NHS Scotland over the past few years. Absolutely. NHS Scotland appears to have moved much further in the interoperability agenda. This is for direct care. For example, it has common laboratory systems, which means that, if you are in Aberdeen or Glasgow, you can see what your patient's blood tests are, as opposed to what happens now between one hospital in north London and another hospital in north London: you would have to ring up and page the doctor at the other hospital to get those results. Scotland is a bit further ahead.

Harry may have a better idea about how that translates into applied research, but certainly for direct care I think NHS Scotland is a bit further ahead than we are in England.

Professor Hemingway: I concur. I think that Scotland has some important lessons about how different parts of the healthcare system can link up data—both for patient care and for research. It is good that Professor Andrew Morris is the director of Health Data Research UK, because of what he brings with him. Previously, he was chief scientist in Scotland, and he brings his understanding of that success to the role.

Q251 Darren Jones: I have a quick question to check I have properly understood the connections between the use of data for direct care and for research purposes. On the question of consent, under the general data protection regulation, there is a distinction between anonymised and pseudonymised data. Anonymised is non-identifiable; pseudonymised is data that have been anonymised but that can be re-identified. In providing direct care, if you rely on anonymised data for research purposes but you re-identify the patient in order to give the care, does that mean that the type of consent ought to be different at the start point?

If you do not know the answer, I suppose that the answer is that we should get an answer.



Dr Hudson: May I clarify the question? There is development of an algorithm and validation of the algorithm, which may be done on a different, anonymised dataset; then there is the application of it in routine clinical use for an individual patient. Those are different datasets.

Q252 **Darren Jones:** Let us suppose I am a patient going into my GP surgery. They do a blood test, measure my blood pressure and check my eyes. The data go on the system. They ask, "Are you happy for this to be used for research purposes?" I say I am, and they say, "That is on an anonymised basis. No one will know that this blood pressure is from you." I want to have direct care, where my data are being re-identified for my purposes. Is it right to say that there is a hard distinction between the use of that data for research compared with direct care, or are there linkages between the use of the data for both purposes?

Chair: Harry, you are nodding wisely.

Professor Hemingway: The former is a fact.

That is where the nub of the invitation of AI, genomics and the vastly expanded data is—that is where it gets to the ethics nub. I do not think we have a clear answer. We have to have the appropriate ethical, governance and consent arrangements that allow you to change your mind over time and allow your direct care providers and those who may be working on research to get back to you, because you may so wish.

On interventions, as you may imagine, all randomised control trials require you to get back to the patient or require access to details to identify the patients to recruit in the first place. It is a big-ticket item.

Q253 **Stephen Metcalfe:** Let us remind ourselves of the potential goal. Dominic, did the streams algorithm you deployed at the Royal Free work?

Dr King: It is not an algorithm that we have deployed; it is a technology.

Q254 **Stephen Metcalfe:** The technology—did it work?

Dr King: It is operationalising the current NHS rules-based, simple, unsophisticated algorithm that has existed for many years. The answer to your question is that, as every technology company should be doing, we are undertaking a rigorous service evaluation of that technology, which is run by some of Harry's colleagues at University College London. We expect to publish that in a peer-reviewed journal in the middle of this year.

I can say that there are some incredibly promising early signs. We receive stories back every day about patients' care being speeded up, some of which have been publicly discussed by the patients themselves. We have nurses at the Royal Free Hospital telling us that it is saving them up to two hours a day compared with their previous ways of working.



HOUSE OF COMMONS

I would prefer to be sitting here saying that we have published a paper in *The Lancet*.

Q255 **Stephen Metcalfe:** Is it still being used?

Dr King: Absolutely—yes.

Q256 **Stephen Metcalfe:** So, it did not stop.

Dr King: No. This is a direct care service. We are providing a critical service now, in keeping with the laboratory systems and the CT scanning software, which delivers a direct care service.

Q257 **Stephen Metcalfe:** Was there a piece of software or piece of technology that you had deployed at the Royal Free that then stopped being used?

Dr King: No. Once we went into live deployment, it continued to be used continuously, and it continues to be used to this second.

Q258 **Stephen Metcalfe:** But, at the moment, the effectiveness of that is still subject to this paper being published. The anecdotal evidence is that it has made a significant difference to patients' lives, which is what this is all about, ultimately.

Dr King: Yes. There are better scientists than me on this panel. We do not like to talk about anecdotes. Nevertheless, I feel very happy when I walk into the ward and speak to the nurses and doctors, and we hear about patients' stories. We are speeding up the time it takes for a senior clinician to see an alert, which, if it is not picked up, leads to patients ending up on dialysis, needing a transplant or, in many thousands of cases every year, dying. As far as I am concerned, speeding up all these processes should translate into better clinical outcomes—but only the evidence will tell us that.

Chair: Is this a quick question, Vicky?

Q259 **Vicky Ford:** I want to ask about the GDPR, which we are about to implement. I remember the huge concern among medical researchers about the early drafts of the GDPR. I remember standing up in the European Parliament saying that it was unworkable for medical research in Europe. Is it now in the right place? Does it give the type of data protections that you have suggested we need? I am asking about medical research—I know that there are concerns about other areas.

Chair: I guess this is for you, Ian.

Dr Hudson: I said earlier that we do not regulate the data; it is the algorithm. With our discussions with the Department of Health, which leads on this, I think, with appropriate safeguards in place, yes. Indeed, the Information Commissioner's Office has reviewed this and has published its conclusions. There are appropriate derogations for research and so on, but there is nothing in the GDPR that will stop the development of this area. But there must be appropriate safeguards in



HOUSE OF COMMONS

place to ensure that rights are protected and appropriate studies are done.

Vicky Ford: Good.

Q260 **Chair:** Before we bring in Bill, I want to ask about a really important point on bias. Dominic, you talked about the importance of getting data from different geographies, ethnic groups, social classes and so forth. How important is that to ensure that we do not build bias into the algorithms that we use in healthcare?

I am also conscious that a certain subset of people have wearables. Is there a risk that we could end up building bias into future decision-making?

Dr King: It is clearly critical that we avoid any bias in these algorithms. If you look at the types of existing algorithm in clinical practice currently, you will see that many of them are fairly unsophisticated and rules based, and they treat every patient as being the same. Whether you are an 18-year-old Caucasian female or an 80-year-old Bangladeshi male with multiple co-morbidities, it applies the same rules-based system. That inevitably leads to false positives and negatives, and all the consequences that follow from that. We at DeepMind think incredibly carefully about that.

The recent announcement that we made on the partnership with Imperial College London, looking at mammography scans, was premised on the fact that we can do this type of study only if we are able to draw data from multiple different geographical sites—ideally, in future, internationally. If it is not a representative data sample, we will inevitably entrench biases into the algorithms, and that will lead to harm. The flip side of all the positives that we are talking about today is that these technologies, if they are not done safely, will lead to patient harm, which we want to avoid.

Q261 **Chair:** Presumably, it should be a national priority—a priority for Government, NHS England and so forth, to have safeguards in place to avoid bias being an issue. Is that right? Harry, you are nodding again, sagely.

Professor Hemingway: Definitely. There are two levels of inequality when it comes to challenges of bias and access. At one level, there are the clinical specialties and places in the country where one has access to data riches. In cardiac surgery, the quality and outcomes have been improved through data for a long time. Is that true for other surgical areas? I suspect not. Some hospitals, such as Birmingham, are great examples of using data to improve patient care and outcomes. That is not the same elsewhere. That is the first level.

The second level is social deprivation, ethnicity, gender and things that the NHS rightly cares about. We know from examples in previous algorithms—for example, for the prediction of cardiovascular disease—



that, if you do not include social factors, you disadvantage the people at highest risk. That is just wrong, and it is unfair. It is absolutely important that we address that.

Dr Hudson: Bias is not a new issue, of course. It has been addressed, for example, by studies that have been done on the clinical practice research datalink for many years. They go through an independent scientific advisory committee review before the studies are done. Bias has always been one of the elements that is part of the review process. Yes, there is a real issue with ensuring that machine learning, algorithms and so on take bias into account and address it, but it has also been an issue for many years in other studies that have been done that do not involve machine learning or algorithms on the CPRD. We have to continue to adapt our approach, but, absolutely, bias has been a risk that has had to be taken into account for many years.

Q262 **Bill Grant:** To what extent is a lack of trust or confidence in healthcare algorithms a problem? Is there a problem with trust and confidence in the use of algorithms?

Dr Hudson: I am not sure that I am the right person to answer that, but I would say that trust is clearly very important for enabling data to be collected and used. On obtaining trust, there must be appropriate safeguards and governance in place, and there must be clear purposes for the use of the data that are available for people. It is key that we get trust and that there is trust across the board from people who know the value of the data in the public health gain—that is the value that I am talking about.

Eleonora Harwich: On top of including all those safeguards—to ensure that we are safeguarding patients' privacy, and proving to medical practitioners that there is value in the use of these algorithms—there is potentially also a question about how the interfaces are designed and how people interact with them. The whole aspect of human-computer interaction is also very important to increase trustworthiness, for both patients and medical practitioners—for example, having a degree of explainability as to how the algorithm reached a certain decision, so that the machine is not totally opaque and shows something of a rationale as to how the decision was made. That definitely helps with the buy-in.

Q263 **Bill Grant:** Is it the patient, the medical professionals, managers or, dare I say it, trusts or boards? Who exhibits the most reluctance or resistance to the adoption of algorithms or to embracing their benefits?

Dr King: If we are going to be successful in this area, we have to overcome the concerns that exist. Some of them have already been picked up. I speak to dozens of clinicians and patients every month, and I think there is general interest and excitement about artificial intelligence in healthcare. There are also some very reasonable questions and concerns.



HOUSE OF COMMONS

Patients care about knowing how their data are being used and who is using their data. We have to do a much better job of telling them how that happens.

As Eleonora has said, there are a couple of common questions that clinicians get. One is on something that we call interpretability. As a clinician, I do not want to see a black-box algorithm spit out a diagnosis to me that I cannot question. We talk about the art of diagnosis—taking many different inputs and coming up with a best-recommended set of actions. I do not think there is a clinician in the world who wants an algorithm that does not allow them to question it. We start from the premise that there needs to be meaningful explanation for the clinician.

Secondly, much like for a pharmaceutical or medical device, clinicians want to have proof that these things work. It is not good enough for me to say, “I think it’s a success,” or, “I think it’s good for you.” We need rigorous, robust clinical evaluations and evidence backing up this work.

The third group of people—organisations, healthcare executives and hospital managers—are very interested, but, as I said earlier, you can see from opening the newspaper that they are fighting fire at the moment. Sometimes, it is difficult for them—reasonably—to think about what the future looks like in five or 10 years’ time when their emergency departments have 10 ambulances outside waiting.

Q264 Bill Grant: You touched on transparency. How do we introduce that transparency to convince these people? What steps could be taken to bring these people on board?

Dr King: From a patient perspective?

Bill Grant: Yes.

Dr King: First, telling patients more saliently how their data are being used. I gave some examples of hospitals being much clearer and research organisations being much clearer about who is processing their data on their behalf. It would be interesting to ask the room: how many of you know the companies that are processing your data at your local hospital? Very few people know that. That is one thing.

Secondly, I said that we are working on a project, the verifiable data audit, which Eleonora also touched on, giving people cryptographic proof that their data are being used in ways such that we can say that we take the highest ethical standards and follow what our legal contracts say. Increasingly, however, we are moving into a world where trust in technology requires the type of cryptographic proof that ledger technologies provide.

Professor Hemingway: Clause 98 of the Data Protection Bill, the right to have knowledge of the reasoning underlying the processing, is useful. That is a positive inclusion, and it comes from the GDPR. There is also a flip side to that: as this conversation expands in society, people



understanding when their data is not being used, which, in my view, does much more harm.

Q265 **Bill Grant:** We have the algorithms and all this data. In parallel with that, we see cyber-attacks or cyber-risks. We think of the military and of banking, but, last May, various health boards throughout the UK, and elsewhere in Europe, I think, fell foul of a cyber-attack. Does that prevent a better roll-out of algorithms and artificial intelligence? Do you view the threat of cyber-attack as something that prevents or reduces the eagerness to progress with artificial intelligence or algorithms?

Dr King: Sanctity of sensitive data is an absolute prerequisite for any use of technology in healthcare. It was very unfortunate that, six months ago, the NHS found itself at such risk from WannaCry. That was not a problem that necessarily affected many other health systems, and that is potentially a sign of problems with chronic underinvestment in technology in the NHS, compared with other industries and other developed health systems.

Q266 **Bill Grant:** So, you reckon that investment could alleviate or mitigate that risk, or perhaps bring comfort.

Dr King: We pride ourselves on having a world-class security infrastructure, which we feel, as much as we can say so, protects patient data. I do not think that anyone should be providing a technological solution, either within the NHS or a third party, that cannot securely manage data.

Q267 **Bill Grant:** You are confident that there are systems that can mitigate or eliminate the risks.

Dr King: Yes—

Q268 **Chair:** Is enough being done across the NHS system to address the concern that that attack highlighted?

Dr King: It was a wake-up call.

Q269 **Chair:** Has there been an adequate response to that?

Dr King: I cannot comment on that.

Chair: Harry? Anyone?

Eleonora Harwich: I do not know how many NHS trusts just updated their systems, because that was what genuinely needed to be done, but it is definitely a wake-up call and a sign that more should be done to ensure that data are stored safely.

Q270 **Vicky Ford:** My questions are about accountability and regulatory frameworks, and what structural systems there should be to deliver accountability for algorithms. I would like to hear your thoughts on that. Should algorithms be subject to audit? Should they be certified before they are used, given that they develop and evolve? Should there be up-



front certification or an ongoing test of an algorithm? Who should be responsible when it fails? Is it Dominic, Harry or Eleonora?²

Professor Hemingway: We will blame Dominic.

Vicky Ford: I guess these are your questions, Ian.

Dr Hudson: Let me start and the others may want to comment. Dealing with static algorithms first, if a company develops an algorithm and it falls within the definition of a medical device, as a number of these will, they have to go through a process of appropriate validation, verification, risk-benefit analysis, appropriate steps in place for monitoring and addressing any safety issues, and so on, that emerge.

They will go through the CE marking process and get it certified. At the moment a lot of these are self-certified, but the new device regulations will make them class II, which means the notified body will review the dossier in due course. As with other products, the developer who goes through getting them CE marked must bear a responsibility for that product. They have to give instructions as to how it is used, and, if used in accordance with those instructions, absolutely, the person whose algorithm it is, who has had it CE marked, must bear a responsibility to ensure the accuracy of that.

Q271 **Vicky Ford:** It sounds to me as if you quite want the CE mark to continue post Brexit.

Dr Hudson: Yes.

Q272 **Vicky Ford:** Can we put that very clearly on the record, please?

Dr Hudson: I think the Government's preferred position is to continue—

Q273 **Vicky Ford:** What is your preferred position? You are giving evidence—not the Government; just your recommendation.

Dr Hudson: In practice, to have a system of continued CE marking would be a very good thing.

Q274 **Vicky Ford:** You are also saying that the international standards do not necessarily give the same breadth at the moment.

Dr Hudson: No, I am not saying that. I am saying a process of review through the notified bodies, which will take into account international standards. The companies will follow international standards or not, but, whatever they do, with the new regulations, they will come to notified

² Note by Eleonora Harwich [Reform]: I think a distinction should be made about the semantics here because I think we are talking about legal liability in case of failure. I think that AI should be following the same procedures than for other medical, where is there a clear recall procedure and liability attached to it. However, I would like to leave the committee members with a question which will be about the role and influence of insurance companies in this debate about assigning liability.



HOUSE OF COMMONS

bodies. Under the existing ones, a lot of these are self-certified. Looking into the future, I think a process like that is necessary, yes.

Q275 **Vicky Ford:** Sometimes I hear people say that we could replace the CE mark with an international, global mark.

Dr Hudson: Certainly, as of day one, we will continue to use the CE mark legislation.

Q276 **Vicky Ford:** What about algorithms that develop and change over time?

Dr Hudson: That is a bit more challenging. It depends on whether it is static today—it is continuing to be reviewed, so there is a new release in a year's time, in which case the new release in a year's time will have to go through the process and so on, but in the meanwhile it is static and then there is a new release—or we are talking about continual learning. I think the continual learning dynamic algorithm is a much more challenging area that we are all thinking about and grappling with, and I do not think anyone has yet worked out the optimal way to ensure that the regulation of these is work in progress, if you like.

Dr King: That is ultimately what we mean by advanced AI—real-time continuous quality improvement. You come into hospital and the algorithm learns from your experience for the next patient. That will clearly present many challenges.

As a general point on regulatory frameworks, the flip side from us, hopefully, having positive clinical impacts is if we do things badly or unsafely or cause harm, and so we absolutely support strong regulatory approaches to any healthcare technology, including artificial intelligence.

On accountability, this is a very interesting question. I am definitely not going to take responsibility for everything, but it is a different conversation if we are talking about an algorithm absolutely replacing a trained clinician. I know some people like to talk about that happening. I just do not see that happening in the foreseeable future at all. I think algorithms will contribute one of many inputs to the final clinical decision-maker. As such, they will remain as a clinical decision aid. That is not to absolve us of any responsibility, as Ian said. It is incumbent on us to make sure that our algorithms work effectively, that we are monitoring and evaluating them. If something goes wrong in that process, we should be absolutely liable for that. But I still see the clinician as being absolutely central to the diagnosis and treatment decisions that patients ultimately have.

Professor Hemingway: I would make two points. First, with regard to the wording of your question, you said “if a company develops an algorithm.” I would hold open the desire and reality that the public sector will develop these algorithms as well. I think that is an intrinsic part of the NHS. I think that the CE marking, whatever its future role, is a small part of the piece. The patient wants to know, “Will I benefit in patient outcome terms?” We have a duty to make sure we translate into clinical



practice things that work and resist those things that do not work. Once things are in clinical practice that do not work, they are hard to extirpate.

Imagine this was a drug. I appreciate there are many differences. Of course, drugs do not change over time, for example. But, if this was a drug, we have an established regulatory framework—all over the world actually—for bringing new drugs to market, if we need to bring new drugs to market. We have a way of dealing with the bias. Every single publication on AI and healthcare that I have ever read is positive. Guess what—that, probably, is not the truth. So, we have the possibility in certain situations of doing randomised controlled clinical trials of some algorithms to say, “Does it shift things that we care about—patient outcomes?”

Q277 **Chair:** You are saying that the same type of rigour should be applied to algorithms as to drugs.

Professor Hemingway: I think that is a good starting point. I absolutely accept that there are differences and it can be a complex intervention. For example, if you imagine a whole patient, which I think is a good thing to imagine, you can imagine many algorithms being used to manage a whole patient, but you could still randomise those patients to many algorithms—

Q278 **Vicky Ford:** But, Harry, where it is considered a medical device, it will have gone through that sort of trial process.

Professor Hemingway: I will let Ian comment, but devices fly under a different radar. Is that fair to say, Ian?

Dr Hudson: There is a different regulatory regime. The answer, though, is that it has to be a proportionate regulatory regime that takes into account the risk and the need, clearly, for patient protection. It must be fit for purpose. I would not say necessarily that the sort of programme you have to put in place for the latest gene therapy product is exactly the same as that for an algorithm, but there must be an appropriate regulatory regime in place that provides everyone with the assurance that is necessary.

Q279 **Vicky Ford:** I think we have very clearly got the message that patients need to be informed and part of the discussion—as well as the medical professional.

Lastly, is there more of a role that should be played by ethics boards? You have said that a management team is dealing with the ambulance queues. Is there more that one could do with ethics boards, their training and so on, to embed them in some of this discussion?

Eleonora Harwich: I think so. The work that DeepMind is doing with having an ethics board and—is it called your review board?

Dr King: We have DeepMind Ethics & Society, which is effectively a think-tank.



Eleonora Harwich: Your research centre within—

Dr King: Then we have a team of independent reviewers.

Eleonora Harwich: That is what I meant. That is definitely an example that other companies should follow, because it is very important to embed ethics within and make sure that all these questions are—

Q280 **Vicky Ford:** But, in the hospital scenario, should the hospital ethics board be involved in this discussion?

Eleonora Harwich: Yes, I think so.

Q281 **Vicky Ford:** Are they adequately trained?

Dr King: In the UK, again, as you know, we are very lucky to have these kinds of very long-established research, governance and ethics boards, which approve and govern the types of health research that we do. Having been to many of those boards, when you introduce new concepts such as laparoscopic surgery, you see that there is a requirement for education and training. I imagine the same will be true as we start bringing some of these clinical implementation studies to ethics boards. It is very important for organisations such as ours, the Wellcome Trust, Farr and various others, to provide the education, support and training so that they can take a reasoned, educated decision.

Chair: Thank you. Carol?

Q282 **Carol Monaghan:** Could I come back to the regulatory issue? Dominic, you said you would be a strong advocate of a robust regulatory framework. Ian, you said it had to be reasonable. I think there is a difficulty here. We have had submissions on both sides. Is there, I suppose, potential that by over-regulating the algorithms we stifle innovation?

Dr Hudson: I think we should also recognise that there are different sorts of regulation, but my comments are particularly around the medical device regulation, where we want a proportionate regulatory regime that, as you say, enables innovation to take place at the same time as protecting subjects. There are the other types of regulation—the GDPR and so on—that we have been talking about earlier that protect, perhaps more, the data.

We must make sure that the regulatory framework both protects the individuals and that the medical device—the algorithm here—does what it says and delivers the benefits that it is being developed for, and at the same time does not cause barriers to it being developed. I agree with you: we must make sure that it is a risk-based, pragmatic regulation that allows that without blocking innovation.

Q283 **Carol Monaghan:** Should best regulation be generic or should it be specific within sectors? For example, does health need its own regulation of algorithms that are being used?



Dr Hudson: I think it is the medical device regulations that I am talking about and the application of those, and ensuring there is guidance for that. I think it has to be at that sort of level. Within the aspects of health that we deal with—medicines and devices—the regulations will be quite different if it is a biosimilar versus a gene therapy product versus a new generic. What you require to do to demonstrate the safety, quality and efficacy will be quite different in those circumstances. Similarly, whether it is an algorithm or one of 500,000 different medical devices that there are now—in fact, it might be a walking stick or an implantable defibrillator—the requirements are really quite different. You have to channel the requirements to the circumstances of the particular algorithm, device and so on.

Eleonora Harwich: I guess this comes back to some of the questions you have been asking throughout the inquiry of whether there should be a general, mega-big algorithm regulator. I do not think that should be the case. Some of the principles applied across different areas could be general, as in we are doing AI for good or we are supporting human flourishing and so on, but I guess the regulations should be specific per sector and build on the existing bodies.

Q284 **Carol Monaghan:** Research Councils UK has said that there has to be an oversight of agreements between public and private sectors. Would you agree with this, or should trusts be free to make their own deals with private companies? Dominic, probably you are the best person to answer.

Dr King: I think a combination of both. If you look at direct care, most trusts enter into specific agreements with individual providers of software and technology. Then it has been very helpful. The reason why we have digital imaging across the UK, which has been incredibly helpful for clinicians and patients, is because national approaches were taken. Similarly, with research, the gold standard for evaluating any type of new treatment or a diagnostic is usually an international, multi-site, randomised controlled trial. If we are going to be successful, research requires data to come from many different organisations. That needs a more joined-up, central approach, alongside some specific projects that would be absolutely fine to run locally.

Q285 **Carol Monaghan:** I suppose the concerns centre around this idea that there could be hundreds or thousands of private companies out there ready to harvest data. They might be doing it for all the best of reasons, but does there need to be some oversight of how these agreements are put in place?

Dr King: It is important to say there are already hundreds of companies working both in direct care and in undertaking research, as well as, as Harry has rightly pointed out, lots of academic groups that are right at the cutting edge of this field too. I get a sense at least that the current processes and regulatory frameworks we have in place are very strong, particularly compared with other countries, but they need to adapt and



change as science and technology move on, so there may well be a requirement or a need for more central oversight.

Q286 Carol Monaghan: Can I move to the testing of these algorithms? I think patient confidence was touched on by Bill already. Do you think it would increase if the results of any algorithm testing were made available to the public? Do you think that would be a useful thing to do? Should these results be made available to the public, and should the NHS be doing the testing of these algorithms?

Professor Hemingway: I would say yes, yes, and more. You would start with: which algorithms have been developed and why? I think there is a reasonable case to say that you start with areas of unmet need that patients themselves identify. They may not be the same as unmet need identified by companies, clinicians or academics. So, start with areas of unmet need. Bring patients to identify those and involve them in every step.

Q287 Carol Monaghan: What do you mean by unmet need?

Professor Hemingway: By unmet need, I mean the fact that, despite for drugs for the treatment of cardiovascular disease, people with a heart attack die from cardiovascular disease. That is unmet need. There are many other examples. We need more interventions—they may or may not be drugs—to tackle that.

Dr Hudson: I think the issue of unmet need is always a difficult one in that a lot of medical advances have happened by small increments, and generally the requirement is that the products demonstrate appropriate safety, quality and efficacy—not that “and there is a need” per se, but appropriate safety, quality and efficacy. To me, the question of who should do the testing does not worry me so much as making sure that appropriate safeguards are in place, appropriate information is available and appropriate governance, and so on, whoever is doing the testing. It should all be done above board, with all the appropriate testing.

Q288 Carol Monaghan: So you think the results of testing should be made available.

Dr Hudson: Certainly, the results of all the testing and so on are available to us. If there is an issue in the marketplace, we ask the company for everything and it is available to us. But I was addressing the question of whether the NHS should do the testing or it could be done somewhere else. To me, the most important element is not so much who does it but that all the appropriate governance is in place.

Eleonora Harwich: In terms of making the results available to patients, there is also a question as to how those are expressed, because I am guessing it is not just the abstract of a scientific journal that is really going to improve buy-in or increase the understanding of patients around that. I think there is the whole work around how that information is presented to make sure that it is effective in reassuring patients.



HOUSE OF COMMONS

Chair: Carol, are you finished?

Carol Monaghan: I have one final question, but can Dominic finish his point?

Dr King: I would like to mention a couple of positive things that we have done. Harry made a good point that only positive articles are published. A good hygiene fact for science now is that you publish before you start the research what you are going to do. In open access journals, we publish exactly the type of project that we are going to run so that someone can come back to us in two years and say, "You've done this study. Why haven't you published it?" Whether the results are positive or negative, this is a real driver now—to publish.

Q289 **Chair:** So you believe in full publication.

Dr King: Absolutely. DeepMind has published well over 150 research peer-reviewed papers in the last two years, including four nature papers. That publishing is at the heart of what we do, and I think we should be open and transparent whether it is positive or negative results.

Q290 **Carol Monaghan:** I have a very final question about the Data Protection Bill that is currently making its way through Parliament. Does this do enough to ensure that algorithms are properly tested before they are used?

Professor Hemingway: Does it do anything?

Dr Hudson: I think it covers the data protection element of the medical device regulations, if you like, for the algorithms that fall within the definition of a medical device. That would define the testing that is necessary to get the CE marking.

Chair: Martin, do you have a question?

Q291 **Martin Whitfield:** A very quick question. Do you see a conflict between the transparency of publishing the results and the need to take the public and the NHS with you, and the protection of the commercial interest in respect of the algorithms? Do you see any concern in that conflict?

Dr King: DeepMind has published some of its most famous algorithms: AlphaGo, the algorithm that beat the world champion in the ancient Chinese game of Go; WaveNet, which is a text-to-speech algorithm and the leading algorithm of its kind, which many of our competitors now use. That was meant to stimulate the research community and show the progress we are making. I do not think we have a red line on saying that we will not publish code. There are very small numbers of people who understand—I certainly don't—the hundreds of thousands of lines of code that make up our algorithm. In the healthcare context, clinicians, and ultimately patients, are more interested in a meaningful explanation of how it works rather than mathematical equations.

In terms of the question on commercial sensitivity, there are issues there, but the more fundamental question from a clinical safety



perspective is that publishing algorithms and releasing them, or allowing other people to release them into the wild, does not necessarily consider the careful safety mechanisms and monitoring mechanisms that we are putting in place to make sure that they are safely deployed. This is so much more than an algorithm. It is about safe implementations and effective monitoring. I am not sure it is particularly useful just to release them.

Q292 **Martin Whitfield:** The algorithm sits within a process.

Dr King: Absolutely—a sociotechnical approach to an unmet need of a problem identified by clinicians and patients. I am not sure how helpful just publishing an algorithm is.

Q293 **Chair:** There is one final question from me. An application of algorithms that has been in the news a bit is the development of online GP services, Babylon being one of them, but there are others. I suppose particularly for Harry and Eleonora, there are different views expressed in the media, the Royal College of GPs expressing real concerns about risk; others are saying that this could be a significant disruptive technology that could seek to address the workforce constraints that we are facing across the system. Do either of you, Harry or Eleonora, have a view, and others as well?

Eleonora Harwich: I have not looked deeply into the Babylon case, but I remember coming across a few blogs in the *BMJ*. One thing that would concern me is the small disclaimer that, basically, if you have chronic health conditions, it strongly advises that you do not use GP at Hand. I do not know if they are going to change that once they potentially scale a little bit more or trial those things a bit more, but I guess there is something to be said around the type of discrimination that that might enact, because, if it is just used by people who are very healthy or who love to track their vital signs and so on, I am not 100% sure how that really is the most important public health improvement.

Professor Hemingway: I would make two points: first, rigorous evaluation. Certainly, evaluation is under way. It remains to be seen how rigorous that is. That is extremely important. The second point is that several GPs have said to me that this can lead to what we do best: the things that computers cannot do. Empathy, judgment and wisdom we find very challenging, and that has an ongoing and changing role in general practice.

Chair: Are there any other comments at all? No. Thank you all very much indeed. It has been an absolutely fascinating session. We appreciate your time.