



Communications and Digital Committee

Corrected oral evidence: Digital exclusion and the cost of living

Tuesday 21 March 2023

3.55 pm

[Watch the meeting](#)

Members present: Baroness Stowell of Beeston (The Chair); Baroness Featherstone; Lord Foster of Bath; Baroness Fraser of Craigmaddie; Lord Griffiths of Burry Port; Lord Hall of Birkenhead; Baroness Harding of Winscombe; Baroness Healy of Primrose Hill; Lord Kamall; The Lord Bishop of Leeds; Lord Lipsey; Baroness Wheatcroft; Lord Young of Norwood Green.

Evidence Session No. 8

Heard in Public

Questions 69 - 73

Witnesses

I: Professor Helen Margetts, Director of the Public Policy Programme, The Alan Turing Institute; Dr Adrian Weller MBE, Director of Research in Machine Learning, University of Cambridge.

USE OF THE TRANSCRIPT

This is a corrected transcript of evidence taken in public and webcast on www.parliamentlive.tv.

Examination of witnesses

Professor Helen Margetts and Dr Adrian Weller.

Q69 **The Chair:** This is our second panel today as part of our inquiry into digital exclusion. Our previous panel was very much about what is happening right now and this next panel is going to help us understand how things may develop in the future. Both witnesses are experts in AI and I ask them to please introduce themselves and the organisation. You are both very much associated with the Alan Turing Institute.

Professor Helen Margetts: I am professor of society and the internet at the University of Oxford, in the Oxford Internet Institute, and I also direct the public policy programme at the Alan Turing Institute, the national institute for data science and AI, where I and quite a large team of people work to try to help public agencies, government departments, local governments and international organisations make the most out of the latest generation of data-powered technologies.

Dr Adrian Weller: I am a director of research in machine learning at the University of Cambridge. I also head safe and ethical AI at The Alan Turing Institute, so Helen and I are colleagues there, and I am programme director for trust and society at the Leverhulme Centre for the Future of Intelligence.

The Chair: As I say, what we are looking at in this session is how AI and algorithms will affect the digitisation of some services.

Q70 **Lord Lipsey:** We all have a tendency, I think, to think of algorithms as something the big tech companies use to direct their wicked ways at us, but I am interested more to hear about the use of algorithms by Governments. How important are they and is this likely to expand hugely in the following years, or has it reached a peak? I would love to hear from both of you.

Professor Helen Margetts: There is one sense in which Governments have used algorithms for a long time, because computer systems first entered government in the 1950s and 1960s and they were based on code: software code is algorithms. In fact, I wrote my PhD about that sort of system back in the 1990s. But the kind of systems that attract so much attention today are algorithms that are powered by data. So, they are based on and driven by data, which the early systems were not, and that makes a number of differences which I am sure we will talk about.

Government departments and agencies do use machine-learning algorithms. Even my programme at the Turing has worked with more than 100 public sector agencies on various aspects. To give a number is very hard: there has been no comprehensive survey of the use of data science and AI in government, really, and of course it is a really difficult thing to measure.

What do we mean here? It could be a passport gate using facial-recognition technology, which is a kind of predictive algorithm, or it could

be some sort of data modelling to understand the implications of different policies—there is such a huge range of things it could be, I could not give a number. A very comprehensive study was carried out in the US, two or three years ago now, by a Stanford group and it found that, even by that time, about half the agencies and departments surveyed were using these technologies in the US. I would say, depending on how you measured it, that that might be true in the UK now.

There is an important distinction to make, though, between central and local government. Central government, as one might expect, tends to be rather more conservative, with a small “c”, and does not rush to be the first to do this kind of thing: remember that technology projects do not have a brilliant reputation in British Governments. Local government has almost been more keen and has probably used these technologies faster than central government. There is an important point to note, which is that local government is being very heavily marketed these technologies by various companies. That is something we might touch on later, because it can be problematic if local government does not know what it is doing and is using these as a kind of black-box technology. That is an important point to make.

One other point: I mentioned legacy systems. These are a bit of a constrainer on using these kinds of technologies, because the kind of systems that I wrote my thesis about did not generate data. Government, particularly central government, does not have a history of using transactional data to do things, so there is a bit of a cultural change that has to go on there. There is a lot more interest in these technologies than there was in earlier generations of technology in government; government is interested in these technologies. When I wrote my thesis, nobody was interested in it. Academia was not interested, you guys would not have been—no one was interested, really, particularly in government. There is so much interest now, and the use of machine learning for public services is mentioned in the Government’s AI strategy, the national digital strategy, the data strategy and so on. There is a lot of interest, which is positive.

Dr Adrian Weller: I would just mention that, on this question and some of the other answers that I may provide, I have been assisted by some of the good people at the Centre for Data Ethics and Innovation, who shared some information and helpfully provided a few numbers. I served on its advisory board. I will take responsibility for any mistakes, but I believe that everything I have to say is correct.

First, I agree with Helen. To my knowledge, there are no centrally held data on exactly what has gone on, but certainly the impression is that it has been increasing. I understand that there was an investigation by the *Guardian* in 2019 that showed that 140 of 408 councils in the UK were using privately developed algorithmic risk-assessment tools, particularly to determine eligibility for benefits and to calculate entitlements.

The Data Justice Lab showed in 2018 that 53 out of 96 local authorities and about one-quarter of police authorities were using algorithms for

prediction, risk assessment and assistance in decision-making. I imagine that we might talk more about those sorts of uses a bit later.

Lots of central government departments are exploring or currently using these sorts of approaches, including the DWP, the MoJ, the Home Office and HMRC. Although much of the current work is exploratory, these different government departments are not always as well connected as one might like to do this optimally.

As Helen mentioned, there are many general tools embedded in all sorts of public service organisations—even things such as email spam detection—and that is growing.

Lord Lipsey: When we look at digital exclusion, we find that some of it is that people are terrified of going digital. From the evidence we have seen, some of that terror is because they fear the use that Governments might make of it. Do you think that will inhibit the growth of this kind of technique in government?

Dr Adrian Weller: It might. Can I take a couple of minutes to explain? A topic which we might get to next but I might start to talk about now, if it is okay, is algorithmic bias and the issues around that. Is it okay if I speak a bit about that now, and connect it, or shall I hold off?

The Chair: If you want to touch on it now to explain something that is fine, but bear in mind that we are coming on to it.

Dr Adrian Weller: I will answer very briefly then, anticipating that coming up later. Yes, there is some reason why reasonable people might have concern about providing their data. That causes the potential for unhelpful cycles where people are worried about providing that data, but not providing it can lead to situations that make things worse for them, which can make them more worried. That is a concern.

Professor Helen Margetts: That can be a concern with any system where you give data. The CDEI has done a big survey on that, and we have just done one at the Alan Turing Institute, which has not yet been released, unfortunately, but will be soon.

However, it is difficult to tell, because if you ask people whether they are worried about something, they are inclined to say that they are. That was one of the things about the CDEI survey. It asked people whether they were worried about giving data to big tech companies, for example, and they said that they were. But people do give their data to big tech companies, so they do it even if they are worried about it.

In the government context, of course, for lots of things, such as if you want to get universal credit, you have to give a certain amount of data; you have no choice. Whether that stops people going digital I am not so sure.

The Chair: Are you, in simple terms, able to identify a difference between the public and private sectors in this area, in terms of the use of predictive algorithms and the data to develop services? As you look at what is happening in government and what is happening in the private sector, is there anything in particular that you would highlight as

significantly different between the two approaches?

Professor Helen Margetts: Generally, the private sector has been using these kinds of technologies more, and did so earlier, than government, but again, there is a huge heterogeneity in the private sector. Amazon has been using predictive shipping for years. In a sense, the big tech platforms invented or at least drove the conceptualisation of these technologies. There will be lots of businesses, particularly smaller ones, which are not using them at all.

The Chair: So it is primarily the fact that they were there, they experienced that and it is part of their business model.

Professor Helen Margetts: Exactly; it is part of their business model. Big supermarkets, for example, are huge leaders in the use of these technologies.

Dr Adrian Weller: I just add that there are two factors to consider. One is the extent to which people feel they can trust who is getting their data. Levels of trust for private versus public organisations have shifted and they fluctuate. I will not comment on exactly where they stand right now, but overall levels of trust are lower than they have been for quite a while.

Of course, an off-setting factor is the extent to which there is an opportunity for a user to use a bright, shiny product that is helpful for them. Organisations such as Google have done very well by providing people with all kinds of services that they really enjoy using and find very convenient. Initially, people might feel worried about their data being used, or perhaps abused, but they start using the product, they like it and they keep using it. They see their friends using the product and, as nothing bad seems to happen, they get seduced to give more and more.

Q71 **Baroness Fraser of Craigmaddie:** I should probably declare an interest in that I am on the board of the British Library, which houses the Alan Turing Institute, although I have nothing to do with that.

I am interested in who is underrepresented or overrepresented in terms of disadvantaged groups. Thinking about the use of predictive algorithms, to what extent is there a risk that, for people who are either underrepresented or overrepresented, that disadvantage will be compounded as we look into the future?

A lot of the evidence we have had has been about services and service design, but are there any examples from other areas? For example, when I read our excellent briefing, it made me think of people trying to get credit scores. If you do not have a credit card or you have not taken out a loan or done anything, you cannot get a credit score. Are there other examples like that, that service providers could look at it to check that predictive algorithms can be used with confidence? How do we know that we are checking all the right things?

Dr Adrian Weller: Credit is a really good example. Thinking about bank loans—I recognise that this is an overly simple way to think about this—a

bank will be likely to make someone a loan or provide them credit if it can have sufficient confidence that the other person will be able to repay the loan. Putting it simply, if I can gather a lot of data on people like you, that will enable me to have more confidence in my ability to predict whether you will repay your loan.

If you are looking at me, a white male, and you know my postcode and income, that might—I am not saying that it necessarily would—be quite sufficient to give you sufficient confidence that you would be able to make me a loan. If I was from an underrepresented group from a postcode of much more highly varying income levels, you might feel that you would require more data about me as an individual to have the same level of certainty about whether I could repay that loan. That is an unfortunate statistical fact, but it is just the way the statistics work.

You can view it in different ways. One way is that it could present an opportunity to say to that individual, "Look, give me some more information and I will be able to make a proper decision about whether to give you credit." On the other hand, many people would feel that that might be perceived as a sort of additional surveillance of someone from an already underrepresented group, who might already feel suspicious about that sort of thing, so they might not want to provide that information. By not providing the information, they are less likely to get a loan themselves, but it will also compound on their community, because people in that community are not providing information—so everyone there is a bit worse off, if that makes sense. So that is one way in which this can be an unhelpful cycle.

Baroness Fraser of Craigmaddie: Helen, the financial loans example is easy to understand, but it looks slightly terrifying if I turn it into a health service scenario, or something. I come back to this question of how we can be confident that the algorithms that we rely on can work.

Professor Helen Margetts: Absolutely. They are based on risk, so they are not completely sure. What is the probability of someone reoffending, for example? If it is 65%, what does that really mean? That is a real problem.

In some data sources, the digitally excluded are underrepresented, and in some they are overrepresented, so it is quite a complex question. To go to the underrepresented, because I felt that that came up quite a bit in the previous session, it all comes down in the end to this question of data and design. These technologies are designed using data so, if your data is not there, it will not be represented in the design. To take an example from healthcare, oximeters in the Covid pandemic, we know that they were not so good at measuring oxygen for darker-skinned people, so they did not work so well for those people. That leads to a more general point that, wherever you are using a source of data, when there is a chunk of people missing, that technology will not be so good at dealing with people like them, whatever they are like. That is an absolutely key point.

One of the speakers in the previous session talked about her mother and about helping her to use the internet and the digital service. That means

that her mother is not represented in the design of that service. The people designing that service did not know that her mother did not use it—and that applies to all sorts of things. Recently, for example, like millions of other people I wanted to query my extortionate gas bill. I went online, and there was a terrible chat thing—so it was not just about digital exclusion—and then I phoned up and it turned out that the helpline was down. They said that there had been too big a volume of calls and that there was no helpline.

In the end, the only thing that got me noticed was going on to Twitter. You cannot be a professor of internet and society and be digitally excluded, so I have a presence on Twitter, and I made a fuss on there—which they were actually recommending. Then they sorted my problem out. So think about their complaints data. Somebody who was digitally excluded could not complain during that period; they had no way to complain. They are not in the data so, if they run any sort of algorithm on their complaints data, those people are not there, and they do not know what their problems are. That applies to all sorts of things—to the sorts of problems that digitally excluded people might have with schools, hospitals or any aspect of public services; they are not in the data. So that is a really significant thing here. Some of your other speakers were saying, “We take great care to think about these people”, but it is really hard for them to do that because they do not know who or where they are.

Baroness Fraser of Craigmaddie: Are we coming back to the point made in the previous session that human intervention is always going to be necessary and that there always needs to be an alternative channel?

Professor Helen Margetts: Yes, but there are things you can do. I was just painting the bleakest picture there. As you heard earlier, lots of people in local government and private companies are doing things to overcome this problem.

What you need to do is to recognise the problem—as with so many problems. You need to realise that there is a whole chunk of people who might be missing here. For example, most people hate filling in forms but, for the digitally included, it is much easier now, because you can do it online, and quite a lot of attention is paid to the design of those forms so that people like us can do them—that is, people who are digitally included. The digitally excluded probably still have to do them offline. That means that there is a very important point about whether their data gets re-keyed into the system, otherwise any predictive algorithm you use will not have those people in it. If it is, you partially overcome that problem. It is a question of working out how the data gets into the right systems and can be treated in the same way as other people’s data.

Dr Adrian Weller: Very briefly, to spell out what may be an obvious point, when you ask how we know that algorithms are going to work, we should recognise that there are economic incentives for most companies in many situations to get their algorithms to work, because that will lead to them making greater profits. But that will not always be the case. The obvious point is that, if the algorithms work well for most people—

Baroness Fraser of Craigmaddie: That is fine for private companies, which do not care as long as they get the number of customers, but public services need to make sure that they reach the people they need to reach.

Dr Adrian Weller: Yes, sorry, I am completely with you, although I understand why it did not sound like I was. I was just pointing out that there is a good chunk of the space where there are good incentives, but there is a really important chunk, on which you are focused, where there are not necessarily good incentives. I am just highlighting the fact that yes, of course, there typically are not those market incentives to serve the digitally excluded. That is exactly what you are focusing on, which is quite right.

I want to make a less obvious point: if you want companies to start to be able to monitor whether they are dealing effectively with underrepresented groups, again, they are going to have a hard time doing that if they do not have their data. Even to be able to check and monitor that effectively is very challenging.

Baroness Fraser of Craigmaddie: You do not know what you do not know.

Dr Adrian Weller: Yes

Q72 **Lord Griffiths of Burry Port:** This question pretty much follows through on the narrative that we have just been hearing. It is about how you would evaluate the response of the Government and local authorities to the challenges raised by digital exclusion in the context of algorithmic data, analytics and prediction. I am delighted that in the potted biographies that we have for each of you, the word “ethics” appeared. That is a great relief to me—but now we will test you on it.

Out of the Alan Turing experience came some guidance for the FAST track principles—“F” for fairness, “A” for accountability, “S” for sustainability and “T” for transparency. There is also a reference to exactly what you have just said, about the risk that this guidance identifies—that the disadvantaged groups may be underrepresented or overrepresented in data samples, and so on. The very successes at the expense of the underrepresented means that they have become even more underrepresented, in the way things work out.

To address this, the guidance recommends that domain expertise will be crucial to assess the fit between data collected or procured and the underlying population to be modelled. How will that domain expertise be acquired and be made useful?

Professor Helen Margetts: I will speak to this, because one of my team wrote it, very much with my huge approval. It comes back to what we were talking about before, a little, because you need to know what the digitally excluded constituency is in your domain. As I said before, there are all sorts of reasons to be cheerful about the way that the Government are using these technologies and the guidance you mention is the official public guidance on the ethical use of AI. The problem, when

it comes to the digitally excluded, is that you need quite fine-grain knowledge on where you are.

I am leading a piece of work for the British Academy on technology and inequality, which will be published tomorrow—annoyingly—but we can send it. That has brought to light how little we know about who is digitally excluded. There is no central place in which we collect data on the digitally excluded and there should be to make that guidance come true because, if you do not know where they are, you cannot adjust for their being missing.

Lord Griffiths of Burry Port: If I may, the councillor who sat where you are was talking about ways in which his operation accumulates information and about the need for partnerships. If, at ground level, the need for partnerships is felt by all and sundry, would the macro level—the people, like you, who run these things—feel the same need for partnerships and accumulation of information? They are on the ground, in libraries, in neighbourhoods, and so on and so forth. If it has to be collected, is a partnership between you and people like that not an absolute necessity?

Professor Helen Margetts: Manchester is a shining example, as one of you mentioned. It has a great big project on digital exclusion and it is really trying to do something. As somebody else mentioned, this is a big, well-resourced council and there has been a big push in that area.

For this British Academy work we ran lots of working groups of people from across local and central government. The trouble was that there were lots of anecdotes and stories, but there was no hard data. I would absolutely recommend some kind of partnership at the centre that collects the data. You cannot resource the problem if you do not know how big it is. If you have a big initiative to tackle digital exclusion, you will not know if it has worked if you do not have any data. Yes, that is a hugely important point.

Stop me if this has come up in a previous session, but nobody in government owns this issue of digital exclusion. I really suggest that that gets some attention, because somebody needs to own and tackle this; it is too important not to.

Lord Griffiths of Burry Port: But we have discovered that this information works across departments and their interests. With DCMS in general, we never quite know where the buck stops.

Professor Helen Margetts: There is a new department now and I know that Ofcom, for example, collects data and the ONS has a lot of data. There are places that have data, but it is just not brought together. If some attention is paid to that, there might be some wins.

The Chair: That is certainly something that we will pick up with the Minister, in a few weeks' time.

Q73

Lord Kamall: I want to ask a couple of quick questions. The first is: to what extent technological innovation, particularly AI—to be honest, I do especially not like the phrase "AI"—could help to address digital

exclusion? I know that this sounds counterintuitive, but is there a way that AI could be used?

Secondly, I have a quick example. I applied for a card—I will say that it was from Monzo—and I was told that I would receive a reply within 48 hours on the app. I got an email a week later saying that I would receive a reply some time. Finally, about two weeks later, I got a reply saying, “You have not been successful and we don’t have to tell you why”. That is fine for a private sector company, in a market where people have choice, but I really wonder about that Monzo interface. I have had a similar experience with easyJet. I would not want to think about easyJet or Monzo technology, know-how or processes being used for public services, particularly in health or other vital issues.

How do we make sure that there is some sort of intervention, whether human or whatever, given the greater reliance on AI that we are seeing? That is my negative point; how do we tackle that? More positively, how can AI be used to tackle digital exclusion? Is there a way that we could use it?

Dr Adrian Weller: First, I embrace your frustration with the term “AI”, which is plastered over everything when we often mean statistical techniques, spreadsheets or even linear regression.

Lord Kamall: It could be large datasets.

Dr Adrian Weller: Yes, although AI can also be used for small datasets. The term is being used to describe everything. But are there general techniques that can help? We should recognise that there are great opportunities to use new technologies to help in all kinds of ways, although there are also concerns.

On your first question about whether these techniques can help to deal with digital exclusion, I note that the first thing people think about is potentially promising and is something that CDEI is investigating, but I will highlight some of the issues. I mentioned before that you might not have data on excluded groups and therefore there is a risk that they could just be ignored. If one tries to do better than that, one first thinks of coming up with an algorithm to try to predict the characteristics of these people and then deal with those. There are ways in which that can be useful, but there are also ways in which it can go badly wrong. As you might expect, you are making up data based on incomplete data that you already have, so you might exacerbate the problems. You have to be very careful doing that. I would not say that it is a complete non-starter, but you have to be extremely careful, look very carefully at the specific situation and understand what is going on. But there is some hope that you might be able to help there.

If you have a bit of data, you can use methods to help address issues of bias. Have we gone past that? Can I talk about it now or can I not talk about it yet?

Lord Kamall: Please talk about it; we are not biased against it, so go on.

Dr Adrian Weller: There is an absolutely correct concern that, by deploying algorithmic systems to make decisions about people, there is a risk of bias that comes from all different kinds of sources. I am not sure that we will have time to really get into that, but I am happy to, if we have time.

Lord Kamall: If you could send that to us as a supplementary in writing, it would be helpful.

Dr Adrian Weller: I am happy to write back, if you could send me specific questions, because otherwise there will be books to fill. One point to keep in mind is that we, as humans, are not devoid of bias—explicit or often implicit. While algorithmic systems are far from perfect and have all kinds of problems, we are developing a toolkit to help us measure and mitigate the bias in these algorithms. To put it bluntly, that toolkit is likely to work better on algorithms than on people. It is easier to adjust these algorithms and understand what is going on in them than it is with people. Even though those algorithms are far from perfect, there is hope that, ultimately, they may be better than humans in some settings. There are a lot of caveats around that.

Professor Helen Margetts: The problem that you have is that you need data to train this wonderful technology. As we have just said, there is a massive lack of data in this area although, as Adrian said, there are techniques around that.

A key thing is the ethical guidance that was mentioned. Most ethical guidance will point to the challenges of AI, which we have talked about. One is fairness and bias, one is accountability or knowing who is responsible if things go wrong, and one is transparency—how do you understand why you did not get the Monzo card or, more importantly, why you did not get universal credit?

We wrote a guide on explainability with the Information Commissioner's Office. That included a lot of citizens' juries, where we talked to people about different domains and different contexts, because context is super important here. Basically, people really mind about transparency if they feel that they could have done something to change what happened, such as they could have got universal credit if they had behaved differently. That is when they care about transparency. With some other things, such as cancer treatment, they are more likely to trust the provider.

I will make two points and then stop. Those three things—fairness, accountability and transparency—are all worse if you are digitally excluded. There is no question about that. You are less likely to understand explanations because they rely on some kind of technological understanding, for example. But there is a bit of a positive light in that, as Adrian said, as humans, we are really biased. We know that. We are biased even when we try really hard not to be, such as when we have done a course on not being biased or whatever. These technologies highlight some of that bias for the first time. We have always had this suspicion about elements of the criminal justice system, for example, but we have not had hard data on it. That might help us. At least we are

talking about some questions that we have not talked about for a long time; we have just assumed that hiring systems are always biased, for example.

The data itself, which, after all, just reflects human behaviour, actually highlights some of these problems, or the lack of transparency and the bureaucracy: the letter which said, "The rules say you can't have...", also lacked transparency. It is making us think anew about some of these values of public administration.

Lord Kamall: That is interesting, because we often think of digital exclusion as part of wider exclusion, but you are saying that we could be using these tools to identify exclusion as a result of bias. Thank you very much, that was very helpful.

The Chair: It is a fascinating topic: the bias or otherwise of algorithms and AI. It is something that we can perhaps explore in another context some other time.

This has been fascinating, in particular when you talked about transparency and the context of the frustration that people have with digital services when it is not transparent to them why a decision has been made. When there is no relationship between that and more conventional, traditional customer service, if I can put it like that, then one feels that ought to have some kind of detrimental effect on the organisations or businesses that do not marry those two things up, if not on their reputation. Maybe there will be, ultimately; I do not know. It seems as though this is not the case for new businesses.

Sorry, I always say something at the end to wrap up that triggers people. I am trying to close this down, but would the Lord Bishop like to say something?

The Lord Bishop of Leeds: I was interested in the language you used. What you said about AI and algorithms implies that, at the end of it, someone learns from the data and what they tell us and then makes a judgment. Actually, the algorithms are making the judgments, such as Lord Kamall not getting his card. A human being probably never intervened in that.

The Chair: That is the connection issue.

The Lord Bishop of Leeds: It is a human agency thing over judgment, decision and ethics. There is a bit of a gap for me.

The Chair: Dr Weller, do you want to say something before I finally close this down?

Dr Adrian Weller: First, that is a whole big topic. I would be very happy to talk about it. In fact, there are many settings where the algorithm makes a recommendation to a human and then the human deals with it. That becomes complex because when they have thousands of these things to deal with they might feel that it is very difficult for them to challenge; they do not have time and they get tired. There are lots of issues around that.

Briefly, on transparency, Lord Kamall's example is really good. The explanation he was looking for was why he did not get his card, in the sense of something he could change so that he could get it. That kind of actionable explanation is very helpful, but we need to be a bit cautious. I am all for transparency and explanation. I do a lot of work on it, but we need to be a bit careful about always asking for an explanation.

As a good example from a few years ago, Amazon recommends products to us. You used to be able to click to say, "Why did you recommend this book to me?" It would say something like, "We recommended book A because you bought book B, and your friend bought book C". We know that really, under the hood, Amazon is very likely doing something more complicated than that; it is giving us an abbreviated explanation. I am not saying that that explanation is not right—there is a whole set of possible explanations you can imagine being given that would all be in part true—but Amazon is probably A/B testing these explanations and it has an incentive, in many ways, to show us an explanation that will make us more likely to click on the book and buy it.

That is actually similar to the way that humans give explanations. If you ask a human why they did something, there is evidence that we subconsciously tend to make up an explanation post hoc to fit what the person wants to hear. You need to be careful in always asking for an algorithmic explanation because it might go in a similar direction.

The Chair: The Chair's privilege is to have the final word on this. That is very interesting, and it is important for us to distinguish between the frustration of engaging with a machine when we want more information as to why the machine is doing something versus a situation where, as a human being, we would expect to engage in a negotiation to reach an outcome that is mutually satisfactory or beneficial to both parties involved. Related to this is when the person one is engaging with is so controlled by a machine that they are not allowed to deviate from it—I have been there and am not going to engage in a further discussion on it; I said that this is the Chair's privilege and I will make my final point. As I said, the whole thing about the biases behind algorithms is a big topic that we might come back to one day in more substance. We will have to invite you back to hear more from you.

Thank you again. That was really interesting and I am very pleased that you were able to talk to us about such a technically difficult topic in a very plain-speaking way. It has been genuinely helpful and informative.