**THE MATTER OF**

**THE IMPACT OF THE PROPOSALS WITHIN "DATA: A NEW DIRECTION" ON DISCRIMINATION UNDER THE EQUALITY ACT 2010**

---

**JOINT SECOND OPINION**

---



## Introduction

1. The Legal Education Foundation (**"TLEF"**) is an independent charitable trust which, amongst other matters, seeks to ensure that the law is developed in a way which is consistent with established equalities and human rights frameworks.[1]

2. TLEF have commissioned us to provide a suite of legal opinions in relation to the equalities and human rights implications of using technology to make decisions about people.

[1] More information about TLEF is available here.

3. In our **First Opinion** (dated the 7 September 2019 and publicly available) we mapped out how the use of data by automated decision making systems to make decisions about people can have profound equality implications and explained how the use of technology should be analysed within the Equality Act 2010 ("**EA 2010**").

4. In this **Second Opinion**, we have a more specific focus. Our advice is now sought in relation to a significant public consultation by the Department of Digital, Culture, Media and Sport ("**DCMS**") called Data: a new direction published in September 2021. This document outlines proposals for reform to data protection laws post-Brexit. TLEF has identified that appropriate reform could ensure that the UK's data protection framework *reinforces* and *promotes* the principle of non-discrimination contained in the EA 2010, in such a way as to avoid the risks highlighted in the First Opinion, while also enabling the better use of these new technologies. That would be a significant gain, but such an outcome is not assured. In this Opinion we discuss some of the possibilities and pitfalls arising from the current consultation.

## Summary of conclusions

5. Section A: The Problem - Tools which deploy automated decision making are increasingly making decisions about people, and there is no doubt that discrimination contrary to the EA 2010 can occur when this happens. We identify four serious legal problems arising from this:

   (i) The current scheme of data regulation does not link properly with the regulation of discrimination;

   (ii) There are low levels of understanding that data processing carries a risk of discrimination;

   (iii) There is an increasingly popular (but wrong and dangerous) view that discrimination does not occur where technology is used in a way which conforms to a statistical measure called "outcome fairness" even though this notion of "fairness" can be at odds with the principle of non-discrimination; and

(iv) Relatedly, the current data protection regime does not effectively allow for the identification of discrimination, as properly defined with the EA 2010, when data processing occurs.

6. Section B: Recommendations - Discriminatory decision making isn't just problematic for the individuals involved; it exposes organisations to legal risk. This is both a burden on business and a distraction from its development for the common good. Indeed, the helpful adoption of technology is inhibited if the public stops trusting organisations for fear of discrimination. So we are clear that the DCMS, through reform of the UK's data protection framework, has an opportunity to fix the problems we have identified that are preventing the principle of non-discrimination from being fully realised in the field of data processing. The right reform could secure that -

(i) The regulation of data processing is placed clearly within the existing equalities framework contained in the EA 2010,

(ii) The principle of non-discrimination is properly understood, rather than wrongly conflated with legally irrelevant notions of statistical fairness, and

(iii) The rights contained in the EA 2010 are more effectively enforced through increased transparency and greater judicial understanding of how discrimination can occur when technology is used.

7. Refinements to the existing data protection regime, which link to the government proposals, and which the DCMS should introduce in order to create positive change include:

(iv) Expressly stating in any new data protection legislative framework that processing which leads to discrimination is unlawful;

(v) Removing the erroneous conflation of "outcome fairness" with equality which is outlined within Data: a new direction;

(vi) Ensuring that comprehensive statutory guidance is created which explains the ways in which discrimination through data processing can occur;

(vii) Making it plain that organisations can legitimately process data in order to check for discrimination, that they must do so and make the results public;

(viii) Ensuring that meaningful, personalised information is provided to individuals where their personal data is processed in order to make decisions about them; and

(ix) Providing the resources to ensure that the judiciary is ready and able to address discrimination complaints linked to discriminatory data processing.

8. Within such a framework the beneficial development of these technologies and the realisation of their full potential will be much easier to achieve.

## Section A: The problem

9. The starting point is that new technological tools - artificial intelligence, profiling and automated decision making - are increasingly making decisions about people; sometimes these systems will not even have a "human in the loop". Data: a new direction contains useful definitions of these concepts (section 1.5, paragraph 66). It further explains how data is used to shape and power these technologies (section 1.5, box after paragraph 66).

10. In our First Opinion we showed how discrimination contrary to the EA 2010 can occur by explaining the application of this legislation to technology through the analysis of case studies relating to the Settled Status Scheme and Risk-based verification ("**RBV**") by local authorities to administer housing benefits and council tax benefits. We concluded that the risks were significant and yet there was minimal commentary in the public sphere about the ways in which data processing could lead to discrimination, and we emphasised the need for a more intense debate about this (paragraphs 167 to 173).

11. Since the TLEF published the First Opinion in 2019, matters have progressed and there is a greater acknowledgment of the risks of discrimination when data is processed.

12. Several contributions stand out:

- The Centre for Data Ethics and Innovation ("**CDEI**") published a report in November 2020 called Review into bias in algorithmic decision-making which noted the potential for discriminatory data processing, for example, in online ad targeting in recruitment (Sections 2.3 - 2.4).

- The TUC published Technology Managing People – the legal implications which identified different scenarios in which discrimination might arise in the context of the employment relationship (Chapter 2A).

- The ICO has published more detailed guidance to organisation highlighting the risks of discrimination from data processing such as Six things to consider when using algorithms for employment decisions which includes advice such as "Bias and discrimination are a problem in human decision-making, so it is a problem in AI decision making".

- There has also been actual and threatened litigation, for example, the Home Office abandoned an algorithm which streamed applicants according to their deemed risk after it was labelled "racist" by the Joint Council for the Welfare of Immigrants.

- The risk of discrimination is properly squarely acknowledged by the government within Data: a new direction.

13. However, there remain four pressing problems. In our view being clear about these problems will enable civil society to push for a better and more relevant reform of the UK data protection regime.

### (i) Regulatory silos

14. First, the effective regulation of technology, to ensure that discrimination does not occur, requires an approach that embraces both data protection *and* the principle of non-discrimination, *at the same time*. However, to date, data protection has been placed in one regulatory "box", while equality has been placed in another; this has created a dangerous blind spot.

15. Equality law recognises that stereotyping is frequently unlawful, yet most modern machine learning processes are involved in trying to find, or perfect, stereotypes

by which to make decisions. So the absence of "joined up" thinking is critical. It is really remarkable now that the EA 2010 makes no reference to data processing, and does not *explicitly* say when, where or how discriminatory data processing is unlawful, and there is no statement in the UK GDPR or the Data Protection Act 2018 ("**DPA 2018**") to the effect that it will always be unlawful to process data in a way which contravenes the EA 2010.[2] The conceptual division between data protection and equality law is historic in origin and it is not necessary; most importantly it has led to siloed thinking.

16. In research which we undertook for Equinet, the European[3] network of equality bodies, we noted this was a common problem; we found a marked absence of co-operation between equality bodies and their data protection counterparts, despite the obvious link between data and new forms of discriminatory technology, like artificial intelligence technology (page 116). It was a significant recommendation that this should no longer continue and we are aware that Equinet has adopted this recommendation as policy.

### (ii) Low levels of knowledge

17. Secondly, in our experience, there are low levels of knowledge that processing data to make decisions about human beings can lead to discrimination within the meaning of the EA 2010. Indeed, in recognition of this problem, in late 2020, we were commissioned by the Council of Europe to produce a bespoke training programme for regulators, which has been rolled out in the UK, France and now Spain. Whilst we hope that this contributed to greater levels of knowledge within the regulation community, we believe that the lack of knowledge continues to be very significant.

[2] While there is a prohibition on certain forms of data processing in relation to "significant decisions" which is defined to mean a decision which "produced legal effects concerning the data subject" (section 14, DPA 2018 & Article 22, UK GDPR), and a decision that led to discrimination would *prima facie* fall into this provision, the section does not apply to all data processing. It also does not state clearly and unambiguously that all discriminatory data processing, as understood under the EA 2010, would amount to a significant decision. We address this further in Section B below.

[3] Equinet is not limited to the member states of the European Union but is the network of equality bodies in all Council of Europe countries.

### (iii) False dogma that "outcome fairness" is a proxy for an absence of discrimination

18. Thirdly, a dogma is developing to the effect that provided that artificial intelligence and related tools are "fair" as judged against a statistical measure called "outcome fairness", then no discrimination can occur. This view is outlined in Data: a new direction at Section 1.5.[4] Yet it needs to be said that this is simply wrong in law; urgent action is required to address this legal fallacy and to stop it becoming accepted. If not, both public and private organisations will be lulled into a false belief that they are acting lawfully when in fact they are breaching the EA 2010. Money and time will be wasted as a result; many will be adversely affected. Sooner or later, this will come back to undermine whatever development has taken place, with untold legal and social consequences. The case for avoiding this result is really that simple.

19. The law recognises that what is "fair" is a slippery notion, which can have many different interpretations in ordinary social or political discourse. By contrast the legal principle of non-discrimination and equal treatment is very well-established and now set out in the EA 2010 in relation to all the major social contexts (work, housing, education, goods facilities and services, and the exercise of public powers). Its provisions are both clear, and, the product of a long iteration of legislation going back at least to 1975.[5] Thus section 13 EA 2010, defines *direct discrimination*, as "less favourable treatment" "because of" a protected characteristic; so for instance it is unlawful to recruit a white person in preference to a black person, where they are both suitable for a role, on the grounds of race. [6] Similarly, section 19 EA 2010 defines *indirect discrimination* in such a way that it

[4] It also wrongly suggests that the EA 2010 is concerned with procedural fairness; it is not as it focuses on outcomes namely that people are treated in a way which is consistent with s.13 to s.21 EA 2010 which focuses on the actual treatment of individuals rather than merely procedural matters.

[5] The basic principle is that unless there is an objective justification, people in analogous situations should be treated in precisely the same way regardless of their protected characteristics, and people in non-analogous situations because of their protected characteristics should be treated differently. This is both a principle of rationality (per Lord Hoffmann in *Matadeen and Others v. M.G.C. Pointu and Others (Mauritius)* [1998] UKPC 9) and the basis of all equality law since it was first enunciated by Aristotle, see e.g. Advocate - General Sharpston's Opinion in *Case C-427/06 Bartsch v Bosch und Siemens Hausgeräte (BSH) Alteredsfürsorge GmbH* and *Regina (DA and others) v Secretary of State for Work and Pensions v Shelter Children's Legal Services and others* [2019] UKSC 21.

[6] The only exception to this rule will be where the positive action provisions in the EA 2010 apply.

will be unlawful for an employer to insist on the same rule for everyone (for example working full time in the office) if that places women at a disadvantage due to childcare commitments, unless that practice can be objectively justified. The duty to make *reasonable adjustments* for disabled people contained in sections 20 to 21 EA 2010 can be analysed in a similar way; the duty to make adjustments only arises where an organisation, such an employer, seeks to apply a rule to everyone but which disadvantages disabled people because of the differences between them and non-disabled people.

20. "Outcome fairness" is defined within Data: a new direction to mean "equal outcomes for different demographic groups". A more detailed definition is used in the CDEI's Review into bias in algorithmic decision-making which states that "Outcome fairness is concerned with what decisions are made i.e. measuring average outcomes of a decision-making process and assessing how they compare to an expected baseline" (page 30).

21. "Outcome fairness" may seem attractive and we recognise that it may have its place in some contexts as a political justification,[7] but it is entirely wrong to assume that it is an adequate approximation for the principle of non-discrimination that underpins the EA 2010. For a start, it fails to recognise that individuals in analogous situations must be treated in the same way regardless of their protected characteristics. It assumes that it is enough to have some broad picture of fairness while ignoring the individual.

22. We can explain this further with an example. Suppose a situation in which a recruitment tool is used to identify 10 candidates for a particular role. There are 1000 applicants, 300 are men and 700 are women. "Outcome fairness" might be used to dictate that 30% of people identified as suitable candidates for the role must be men and 70% must be women meaning that the final recommended pool should consist of 3 men and 7 women. However, if only 2 men within the 100 applicants are the most suitable candidates, it means that one man will have advanced to the pool even though he does not meet the right standard. Further, if there were 8 women who were most suitable, one woman would need to be "held back" so that 3 men could be put forward and the "right" statistical outcome achieved. In such a scenario here will be direct sex discrimination as understood within the EA 2010 as one better suited woman will not advance so that a less suitable man can be placed in the pool. Whilst a tool which created this type of outcome might seem

[7] For instance, it may be politically necessary in relation to the distribution of medicines in a pandemic.

"fair" from a purely high level and statistical perspective, in truth it is discriminating because of the actual impact which it has on certain individuals.

23. Interest in "outcome fairness" as a proxy for the principle of non-discrimination appears to have developed in United States where the concept of equality is not identical to that which we have. The US concept developed from its particular history of intentional prejudice against the black community. Thus, the law infers intentional prejudice when a pattern or practice of preferring (for instance) white recruits over black, is demonstrated to a sufficiently high degree. The inference is said to arise when there is an 80% or worse chance that a white person would be chosen as compared to a black person. However, this is neither part of UK nor European anti-discrimination law as the CDEI Bias Review pointed out. Here, there is no need to prove intention to establish discrimination.[8]

24. It is imperative that organisations in the UK understand the provisions of the EA 2010 rather than think that such "outcome fairness" is enough. Whether or not it is for the US, it is not for the UK. A failure to be clear about the difference will mean that organisations will unwittingly discriminate here in the UK. Equally in so far as they operate in Europe, they are also very likely to be breach in European definitions of equality which are the same as the notion used in the UK.[9]

### (iv) Opacity

25. Fourthly, making decisions about human beings by deploying machines which process personal data leads to a lack of transparency and often to the decisions not being readily observable. We call this opacity. Jennifer Cobbe in a 2019 paper[10] identified three types as follows:

- Intentional opacity: This will occur where a system's inner workings are deliberately concealed to protect intellectual property;

[8] For example, see the Supreme Court decision in R (on the application of Coll) (Appellant) v Secretary of State for Justice (Respondent) [2017] UKSC 40 where direct discrimination arose due to the application of rules where there was no human actor who intended for discrimination to occur.

[9] Please note that the European Union is intent on introducing a new legal regime called the AI Regulation which will expressly prohibit the discriminatory use of AI within Europe.

[10] For a more detailed explanation of the ways in which these forms of opacity can arise, and links to Jennifer Cobbe's paper, please see the First Opinion, paragraphs 9 to 13, 127 to 143 and 155 to 162.

- Illiterate opacity: This will occur where a system is only understandable to those who can read and write computer code; and

- Intrinsic opacity: This will occur where the complex process of machine learning means that a human will struggle to understand the decision-making process.

26. We also add a fourth type of opacity which will arise when data sets are deleted due to concerns, driven by the UK GDPR, that data should only be processed for its collected purpose. This means that the data that was processed within an automated tool is swiftly deleted after a decision has been made making it hard to then understand afterwards whether or not the data processing was discriminatory.

27. In short, the outcomes from these automated systems are frequently not transparent, too difficult for even experts to explain and understand in any detail. People may not even know that decisions are being made about them let alone *why* a specific decision has been made about them. This should be contrasted to when humans make decisions about each other. Humans give an explanation for their decisions, and where a discrimination claim is brought, they can be required to attend to give evidence and be cross-examined about their decision and the reasons for it. So, while we recognise that it is common to say "AI can discriminate but so can humans"[11], we need also to be clear that humans can explain themselves in a way in which AI machine learning systems cannot. The ability a court or internal appeal procedure or review to examine the explanation for a human decision, is an important procedural step, when it comes to ensuring that discrimination does not happen. Whether human decisions are good or bad there is a real legal value in the fact that they are capable of being explained in this sense. In practice it is a safeguard of the right to be free of discrimination and to have equal treatment, for *both* putative victims and those persons or undertakings that are alleged to have discriminated. We address later the level of explanation which we believe that automated systems should be required to provide to demonstrate that they are not acting in a discriminatory way (see paragraph 49).

[11] Seem, for example, the ICO document Six things to consider when using algorithms for employment decisions which states "Bias and discrimination are a problem in human decision-making, so it is a problem in AI decision making".

28. We have noted in our First Opinion that individuals *might* identify discrimination through a request for transparency pursuant to Article 15 of the UK GDPR in so far as personal data is being processed. This obliges a Data Controller to explain the categories of personal data being processed and in particular –

> the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject

29. The first point to note is that this right to information is qualified. It only applies where Article 22 applies, i.e. there is fully automated decision making as opposed to where a human makes a decision in conjunction with an automated system. Secondly, some care should be taken in relying on Article 15 (even where Article 22 is engaged) in the UK. Whilst the Amsterdam District Court recently relied on the GDPR in the Ola litigation in favour of the drivers finding that "*useful information about the underlying logic*" of an algorithm must be deployed[12], the guidance on automated decisions given by the UK's ICO suggests that the principle of transparency is fairly weak when it comes to algorithms. The ICO's position is that Article 15 in the UK GDPR does not extend to providing detailed, *individualised* explanations of why a decision has been made such that it could be ascertained whether discrimination is occurring or not. Indeed, the ICO provides the following example of what explaining a decision would look like in practice:

> An on-line retailer uses automated processes to decide whether or not to offer credit terms for purchases. These processes use information about previous purchase history with the same retailer and information held by the credit reference agencies, to provide a credit score for an online buyer.
>
> The retailer explains that the buyer's past behaviour and account transaction history indicates the most appropriate payment mechanism for the individual and the retailer.
>
> Depending upon the score customers may be offered credit terms or have to pay upfront for their purchases.

[12] Please see the excellent blog by various academics in the Netherlands.

30. Further, the existing Article 35 UK GDPR requirement to analyse processing that assesses the risks to the rights and freedoms of data subjects by means of a Data Protection Impact Assessment is of little utility because the only obligation is to provide it to the ICO in certain circumstances and not to publish it generally.[13]

31. Lastly, there is also a mechanism within the existing data protection regime for people to access their data which are known as Subject Access Requests. It is always possible that an individual could make a Subject Access Request and from this learn that data relating to their protected characteristics is being processed which might suggest that discrimination is happening. However, in the absence of a meaningful explanation as to how that data is being used, simply being provided with access to the data that is processed will likely be insufficient to understand whether discrimination has occurred.

32. In summary, no effective mechanism has yet been provided in the UK through the existing data protection framework to enable individuals to readily understand how decisions about them are being taken and yet these decisions carry a risk of discrimination.

## Section B: Recommendations

33. Discriminatory decision making isn't just problematic for the individuals involved. It exposes organisations to legal risk. It will also prevent innovation if the public stops trusting organisations that process data if there is a fear of discrimination.

34. Indeed, surveys show that trust in technologies like AI are now very low in the UK. The CDEI published a Barometer Report on 23 June 2020 finding significant levels of lack of trust in AI. Since then, it has certainly deteriorated because of the controversy concerning the decisions made between Ofqual and the Department of Education during the summer of 2020.[14] The fact that students were being "*marked*" by algorithm rather than human examiners suddenly alerted the public to the fact that life-changing decisions were being made about people "*by machines*"

[13] See the ICO guidance available here and an article by Swee Leng Harris, Data Protection Impact Assessments as Rule of Law Governance Mechanisms June 3, 2019: https://zenodo.org/record/3237865#.XTGSTPJKhQI

[14] The Chair of Ofqual's written statement to the Education Select Committee concerning the award of GCSE, AS, A levels in 2020 is here.

and not humans and for many this was unacceptable. A report of the British Computer Society published on 22 September 2020 after the exam fiasco contained the message:[15]

> …the majority of people do not trust computers to make decisions about any aspect of their lives…

35. The DCMS, through reform of the UK's data protection framework, now has a golden opportunity to ensure that the principle of non-discrimination is placed at the heart of the regulation of data processing.

36. There are refinements to the existing data protection regime, which link to the government proposals, and that the DCMS can introduce in order to create positive change.

### (i) Legislative statement that discriminatory data processing is unlawful

37. The dangerous blind spot created by placing data protection in one regulatory "box" and equality in another which we described in paragraph 14 above can be remedied by amending the DPA 2018 and UK GDPR so that it states unequivocally, and without any exceptions, that data processing which leads to breaches of the EA 2010 is unlawful.

38. As part of the consultation process, the government asks at Q1.5.1, 1.5.2 and 1.5.4 for feedback on how fairness should be defined in a new data protection framework. We are clear that new legislation should recognise that fairness is conceptually different to the principle of non-discrimination and that no data processing can be lawful where it discriminates as understood within the EA 2010. It is simply not enough to say that processing must be fair if it is not also non-discriminatory.

39. A legislative statement to this effect would be different to existing Articles 21 and 22 in the UK GDPR. Those provisions merely create heavily qualified rights to

[15] The British Computer Society's Report is here.

object to data processing and a prohibition on data processing which produces "*legal effects concerning*" an individual or "*similarly significant affects him or her*". Whilst some discriminatory data processing will certainly fall within the ambit of these provisions, we believe that *all* discriminatory data processing as understood by the EA 2010 should be expressly stated to be unlawful without any qualification.

### (ii) Removing the erroneous conflation of fairness with the equality

40. Data: a new direction recognises that there is uncertainty about the scope and substance of "fairness" in the data protection regime (Section 1.5, especially paras 76 to 79). It further recognises that "outcome fairness" in the context of AI within the data protection regime may not be "*feasible or effective*". We agree entirely with that sentiment but go one step further. The inaccurate conflation of fairness with the principle of non-discrimination contained in the EA 2010 as outlined paragraphs 18 to 24 above is dangerous and should be removed entirely from the future discourse. There must be recognition that concepts like "outcome fairness" are false proxies for the principle of non-discrimination and that systems which merely create "outcome fairness" could breach the EA 2010. This recommendation dovetails with our proposal in paragraph 37 above that the new data protection regime must build directly on the EA 2010, using its language and the principle of non-discrimination which it embodies.

41. As part of the consultation process, the government asks at Q1.5.1, 1.5.2 and 1.5.4 for feedback on how fairness should be defined in a new data protection framework. We are clear that new legislation should recognise that fairness is conceptually different to the principle of non-discrimination and that no data processing can be lawful where it discriminates as understood within the EA 2010.

### (iii) Comprehensive statutory guidance

42. To support organisations further understand what is required of them when they process data and how discrimination can occur, there should be comprehensive statutory guidance which is sector specific.

43. As part of the consultation process, the government asks at Q1.5.3 for feedback on what legislative regimes and associated regulators should play a role in substantive assessment of fairness, especially of outcomes, in the AI context. We are clear that statutory guidance produced by the ICO and the Equality and Human Rights Commission ("**EHRC"**) which explained the difference between fairness and non-discrimination and give practical guidance to organisations about how to act lawfully is urgently needed.

### (iv) Legitimacy of data processing in order to detect discrimination

44. The principle of non-discrimination is meaningless without adequate means to identify that discrimination is happening. One way in which the opacity problem identified in paragraphs 25 to 32 can be remedied is by making it plain that organisations can legitimately process data in order to check for discrimination.

45. Accordingly, we support the Government recommendation that organisations should be able to process data where it is necessary for "*monitoring, detecting or correcting bias in relation to develop AI systems*" (Section 1.4, especially paragraphs 60 to 61). In terms of terminology, we believe that any new legislative provision should expressly refer to the EA 2010 so as to reinforce the link between the data protection regime and the principle of non-discrimination and be broader than simply "*AI system*", for example, by referring to "*monitoring, detecting or correcting discrimination when data is processed contrary to the EA 2010*".

46. But this is not the end of the matter, we are also clear that organisations should be obliged to undertake equality assessments of their tools when they are processing data *and* that the results should be made public. Creating a spotlight on inequality is a powerful first step towards its correction as recognised in the employment field with the obligation to report Gender Pay Gaps.

47. We recognise that regulation must be proportionate and appropriately targeted, and as such we suggest that the requirement to undertake assessments and publish them should be subject to a minimum threshold, for example, only where data

processing is deemed "high risk" to the principle of non-discrimination which is the mechanism which is currently used in Article 35 of the UK GDPR.[16]

48. As part of the consultation process, the government asks at Q1.4.1 to 1.1.43 for feedback on whether data should be processed in order to identify discrimination. We are clear that new legislation should recognise that this type of monitoring should take place in terms outlined in paragraph 44 and that where the data processing is "high risk", the monitoring should be compulsory with an obligation to publish.

### (v) A requirement for personalised explanations

49. A further way in which the opacity problem identified in paragraphs 25 to 32 can be remedied is by ensuring that meaningful, personalised information is provided to individuals where their personal data is processed in order to make decisions about them. As we have already explained, no such mechanism currently exists within the data protection regime in the UK. This is not a matter that is foreshadowed in with Data: a new direction.[17] In our view, the Government should really focus on ensuring that meaningful, personalised explanations for decisions are mandatory in the way that we have proposed. In order to ensure that such an obligation was proportionate, it may be appropriate to limit that obligation to only "high risk" uses of technology as explored in paragraph 47 above. However, as we have consistently advised, identifying what is appropriately classed as "high risk" requires joined up thinking as between data and equality regulators.

[16] There is a developing debate in Europe on this topic. The EU rightly recognises the importance of proportionate legislation and proposes a threshold for regulation of "high risk" applications of artificial intelligence as the means of doing so within its proposed AI Regulation. How to define "high risk" appropriately and in a UK context would require further and detailed thought.

[17] Although we note that section 1.5 examines discrimination and 2.3 looks at reform of the Subject Access Request regime.

### (vi) Capacity building

50. Finally, the DCMS should ensure that the judiciary is ready and able to address discrimination complaints linked to discriminatory data processing. Data: a new directions, section 5.6, examines whether reforms are required to the existing complaints regime. However, we consider that bigger point is at play which is that if non-discriminatory data processing is to be avoided, there must be an enforcement of rights by the courts and for this to happen effectively a capability building exercise is needed in terms of judicial knowledge of this area.

51. Some steps are already taking place in relation to this. For instance UNESCO has just published its intention to develop a "'massive open online course" or MOOC called "AI and the Rule of Law: Capacity Building for Judicial Systems".[18] This will take some time to develop and will probably be at some level of generality but similar proposals are also under discussion in the Council of Europe.[19] The key point is that the development of these new data technologies requires new skills from the judiciary if they are to uphold the rule of law. It may well also require new processes as well; we note that the Futures Group of the Civil Justice Council is to engage with some of these issues.[20]

52. We end by reminding readers that there can be no democracy without equality before the law; anything less is not consistent with the rule of law. The huge increase in data and data processing and new technologies beg really significant questions about this kind of equality: how is equality before the law to be provided if one party has all the data and an unintelligible process for using that data, which is in any event guarded by significant intellectual property rights? The fullest answer to that question is outside the scope of this Opinion but its importance cannot be in doubt. Data protection legislation to date has tried to go some way to

[18] Its aim being to 1.) Stimulate a participative dialogue with judicial operators on AI-related innovations in the judicial system and promote knowledge of digital innovations in the justice system; 2) Facilitate knowledge exchange and experience sharing among judicial operators on artificial intelligence, existing norms and standards (hard and soft law) in the field, and its implications for human rights; 3) Highlight existing case studies and best practices that translate ethical principles into practice both in terms of the use of AI in justice systems, and in cases involving AI impacting human rights.

[19] See the work of the Council of Europe's Commission for the Efficiency of Justice here.

[20] See its terms of reference here; note the intention to "… take and encourage a long-term view of the impact of technology on the administration of justice, with emphasis on increasing access to justice and securing the position of the legal system of England and Wales as a global leader."

balancing the rights of data subject and data processor within a regulatory environment. To date, this balance has not been drawn in the right place when it comes to ensuring that the principle of non-discrimination is properly respected. This consultation is an opportunity to ensure that data protection and equality rights are enhanced and certainly not ignored.

**ROBIN ALLEN QC** **DEE MASTERS**

ra@cloisters.com deemasters@cloisters.com

www.cloisters.com

www.ai-lawhub.com

5 November 2021