Public Law Project

Department for Digital, Culture, Media and Sport's consultation

'Data: a new direction'

Public Law Project's Response

Contents

Introduction

1. Public Law Project (PLP) welcomes the opportunity to respond to the Department for Digital, Culture, Media, and Sport's (DCMS) consultation, 'Data: A new direction'. We confine our response to selected questions within Chapters 1, 2, and 4 that lie within our specific expertise.

2. PLP is an independent national legal charity. For 30 years, PLP's mission has been to improve public decision-making and facilitate access to justice, through a mixture of casework, advocacy and research. PLP's vision is a world in which individual rights are respected and public bodies act fairly and lawfully.

3. The focus of our response is on questions of relevance to government use of automated decision-making (ADM) systems, sometimes described as algorithmic decision-making. As an organisation, this is one of our four core focus areas.

4. A note on terminology: by 'algorithm', we mean a set of rules for performing a task or solving a problem. This includes, but is not limited to, computer-operated algorithms. By 'automated decision-making', or 'ADM', we mean decision-making processes that are fully or partially undertaken by a computer-operated algorithm. For the most part, we use the term 'ADM' or 'ADM system' in preference to the term 'algorithm' or 'algorithmic decision-making'. However, we recognise that the term 'algorithm' is often used to mean a computer-operated algorithm. In this document, we occasionally use the term 'algorithm' in this way.

Key points

Fairness in an AI context

- In PLP's view, the courts and their interpretation of all the requirements of the EA 2010 have a valuable role to play in assessments of fairness in the context of ADM.
- Any UK GDPR definition of fairness should not usurp or contravene other established interpretations of fairness, equality, and non-discrimination in the law.
- We endorse the recommendation of Dee Masters and Robin Allen QC that the DPA 2018 and UK GDPR should be amended to state "unequivocally, and without any exceptions, that data processing which leads to breaches of the EA 2010 is unlawful".

Article 22

- Article 22 – or something like it – is essential for ensuring human oversight in public decision-making.
- PLP supports reform of Article 22, clarifying its key terms – especially "a decision based solely on automated processing" – to ensure that it has broad practical application.
- Article 22, properly defined, should prohibit *de facto* solely automated decision-making where, due to automation bias or for any other reason, the human official is merely rubber-stamping a score, rating or categorisation determined by the computer. It should require meaningful human oversight, rather than a token gesture.
- We would strongly resist a narrow definition of "legal" or "similarly significant" effects.
- PLP strongly disagrees with the proposal to remove Article 22 of the UK GDPR.

Data Protection Impact Assessments

- We agree with Swee Leng Harris that, "transparency and accountability would be enhanced if DPIAs by government were published by default, so that there was a publicly accessible database of government DPIAs." Mandatory publishing of

government DPIAs on a centralised, searchable and publicly accessible database would help to uphold rule of law standards.

- In our view, a requirement for routine DPIAs, or a similar form of evaluation, throughout the period of deployment of an ADM system ought to be considered. A measure of this kind may help to identify unanticipated rights-violations.
- PLP disagrees with the proposal to remove the requirement for organisations to undertake DPIAs. We consider that more should be done to enhance the utility of DPIAs. They should not be done away with.

Data subject access requests

- PLP considers that reduction in the cost of responding to data subject access requests is not a sufficient justification for curtailing the rights of data subjects. We strongly disagree with the proposals to introduce a cost limit and amend the threshold for response.
- PLP considers that the case against introducing a fee for processing data subject access requests is stronger than the case for it. We strongly disagree with the proposal to introduce a fee.
- We disagree with the proposals relating to data subject access requests regardless of whether compulsory transparency reporting is introduced.

Government algorithms and compulsory transparency reporting

- In principle, PLP supports compulsory transparency reporting. This is subject to two important caveats: first, this must be meaningful transparency; and, second, transparency alone is not enough. Subject to these two caveats, we consider that compulsory transparency reporting has the potential to help government earn public trust in their use of data.

- PLP suggests that compulsory transparency reporting requirements should apply at minimum to high risk systems, broadly and flexibly defined. The possibility of including a non-exhaustive list of high risk systems could be considered.
- PLP suggests that DCMS should consider whether other parties providing or using high risk ADM systems should be subject to compulsory transparency reporting requirements, too.
- In order to achieve meaningful transparency, PLP suggests that the following types of information should be considered for inclusion in compulsory transparency reporting requirements: details of the provider and user(s) of the ADM system; identifying details; purposes; status, including period of deployment; any training data, methodologies, and techniques; DPIAs; Equality Impact Assessments (EIAs); an explanation of how the system works; and an executable version of the system.
- PLP recognises that some exemptions may be necessary.
- An obvious model for any such exemptions would be the FOIA. However, PLP has some reservations about this approach.
- We do not think exemptions should apply to the disclosure of high-level information, including the fact that there is an ADM system in use in a particular context. If more detailed information about a particular system is withheld on the basis of an exemption, there should be readily accessible avenues for challenging this, with the possibility of review by an independent regulator such as the ICO.

Chapter one – Reducing barriers to responsible innovation

Fairness in an AI context

5. Article 5(1)(a) of the UK General Data Protection Regulation (UK GDPR) requires that personal data must be processed lawfully, fairly, and in a transparent manner. The starting point is that fairness is a stand-alone principle and obligations of fairness under Article 5 should be treated as conceptually distinct and independent from obligations of transparency.[1] Transparency and fairness should not be conflated or absorbed into one another.

6. We agree with the Information Commissioner's Office (ICO) that while fairness historically focused on transparency:

> [F]airness now also ensures that the way in which people's data is processed does not lead to unfair or unjustified impacts on their lives. This is important because simply telling someone what you are doing with their data does not, by itself, make it fair [emphasis added].[2]

7. Unfairness can exist where ADM technology systematically reproduces existing biases or inequalities in society or where it is trained mostly on data from one group of people and then applied to another group where it produces worse outcomes.

8. The consultation asks: *Q1.5.3. What legislative regimes and associated regulators should play a role in substantive assessments of fairness, especially of outcomes, in the AI context*?

---

[1] Orla Lynskey in Jorrit J Ripma (ed) The New EU Data Protection regime: Setting Global Standards for the Right to personal Data Protection Vol. 2 at 38.

[2] ICO, Response to DCMS consultation "Data: A new direction", 6 October 2021, available at https://ico.org.uk/media/about-the-ico/consultation-responses/4018588/dcms-consultation-response-20211006.pdf.

9. The Public Sector Equality Duty (PSED), set out in section 149 of the Equality Act 2010 (EA 2010) and elucidated by the Court of Appeal in *Bridges*,[3] is important in ensuring fairness for those subject to ADM.[4] As is well-known, the *Bridges* litigation concerned the South Wales Police's (SWP) use of facial recognition technology. Before the High Court, there was evidence that, due to imbalances in the representation of different groups in the training data, such technologies can be less accurate when it comes to recognising the faces of BAME people and women.[5] The Court of Appeal found that the PSED imposes a duty on public bodies using ADM systems to take reasonable steps to gather relevant information about the impact of the system on people with protected characteristics.[6] Public bodies must give advance consideration to issues of discrimination before making a decision to use a technology. They must make their own enquiries, rather than uncritically relying on the assurances of a third-party contractor.

10. It is clear, then, that the PSED encompasses what DCMS has referred to as 'outcome fairness'. The PSED required the SWP to do "everything reasonable which could be done… in order to make sure that the software used does not have a racial or gender bias";[7] in other words, to ensure that the software was no worse at recognising the faces of BAME people or women. It is arguable that a similar duty should be extended to private parties using ADM systems.

11. However, fairness under the EA 2010 is not limited to 'outcome fairness'. The EA 2010 is not solely concerned with fairness on a statistical level but with the treatment of individuals. Dee Masters and Robin Allen QC give the following example:

---

[3] *R (Bridges) v South Wales Police* [2020] EWCA Civ 1058.

[4] Jack Maxwell and Joe Tomlinson (2020) Proving algorithmic discrimination in government decision-making, Oxford University Commonwealth Law Journal, 20:2, 352-360.

[5] See the expert report of Dr Anil Jain, available at https://www.libertyhumanrights.org.uk/wp-content/uploads/2020/02/First-Expert-Report-from-Dr-Anil-Jain.pdf.

[6] *R (Bridges) v South Wales Police* [2020] EWCA Civ 1058.

[7] Ibid, at [201].

> 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… 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.[8]

12. This example, judged by the standard of 'outcome fairness' alone, may seem unproblematic. But, under section 13 of the EA 2010, the woman who is "held back" is subject to direct discrimination on the basis of sex. We should not lose sight of the unfairness of direct discrimination, simply because we are in an ADM context. In PLP's view, the courts and their interpretation of all the requirements of the EA 2010 have a valuable role to play in assessments of fairness in the context of ADM.

13. The consultation asks: *Q1.5.4. To what extent do you agree that the development of a substantive concept of outcome fairness in the data protection regime - that is independent of or supplementary to the operation of other legislation regulating areas within the ambit of fairness - poses risks?*

14. We acknowledge that certainty in the law is important and there may be utility in a specific concept of fairness for the purposes of UK GDPR. However, fairness in an ADM context should not be limited to 'outcome fairness'. Further, any attempt to define fairness in the UK GDPR should reflect the fact that law does not exist in a vacuum. Any UK GDPR definition of fairness should not usurp or contravene other established interpretations of fairness, equality, and non-discrimination in the law.

[8] Robin Allen QC and Dee Masters, 'Joint second opinion in the matter of the impact of the proposals within "Data: a new direction" on discrimination under the Equality Act 2010' (5 November 2021), available at https://research.thelegaleducationfoundation.org/wp-content/uploads/2021/11/TLEF-Second-Opinion-5-November-2021.pdf.

15. We agree with Dee Masters and Robin Allen QC that the Data Protection Act 2018 (DPA 2018) and UK GDPR should not be "siloed" off from the EA 2010. We endorse their recommendation that the DPA 2018 and UK GDPR should be amended to state "unequivocally, and without any exceptions, that data processing which leads to breaches of the EA 2010 is unlawful".[9]

Article 22 GDPR

16. Article 22 of the UK GDPR provides that a data subject shall have the right "not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her."

17. Article 22 – or something like it – is essential for ensuring human oversight in public decision-making. It also helps to ensure transparency: under Article 22, in combination with Articles 12 and 13, data subjects can get information to find out if they are being subjected to solely automated decision-making.

18. The consultation asks: *Q1.5.14. To what extent do you agree with what the government is considering in relation to clarifying the limits and scope of what constitutes 'a decision based solely on automated processing' and 'produc[ing] legal effects concerning [a person] or similarly significant effects'?*

19. In our experience, it appears that Article 22 is capable of placing a meaningful limit on the deployment of ADM systems. For example, in response to PLP's request for information about Her Majesty's Prison and Probation Service's (HMPSS) use of algorithmic decision-making, a representative said:

> I hope it is helpful if I explain that HMPPS does not, and will not use computer generated algorithms to automate or replace human decisions in HMPPS that have any significant impact on staff, people in prison or on probation. We may seek to

[9] Ibid.

automate low-level administrative decisions but this will always be deployed with human oversight and stringent quality assurance measures. We do not have any examples where 'a computer automatically performs all of the decision-making process without a human's direct involvement'.[10]

20. This suggests that at least some government departments may be interpreting Article 22 broadly, so as to meaningfully restrict the role of automation in decision-making.

21. Nonetheless, we agree with Lilian Edwards et al that Article 22 as currently drafted could be interpreted much more narrowly, such that many if not most public ADM systems are excluded from its scope. They give the following example: Ofqual argued that the allocation of GCSE and A-level grades using their standardisation model was not a solely automated decision because some of the inputs were generated by teachers i.e. by human beings. However, virtually all real-world ADM systems will use human-generated inputs. So, on Ofqual's interpretation, the scope of Article 22 would be very limited. [11] Hence, PLP would support reform of Article 22, clarifying its key terms – especially "a decision based solely on automated processing" – to ensure that it has broad practical application.

22. There are two main types of ADM system. In a rule-based system, the rules have been written by the person designing the system. In a statistical system, instructions are drawn from patterns in historical data. This can be done by a human statistician, or it can be done automatically through machine learning.

23. Another set of variables is the way the automated system fits in to the overall decision-making process. On one end of the spectrum, a decision may be fully automated. By "fully

[10] The relevant freedom of information request and the response are available at https://www.whatdotheyknow.com/request/information_about_use_of_algorit#incoming-1673799.

[11] Lilian Edwards, Rebecca Williams, and Reuben Binns, 'Legal and regulatory frameworks governing the use of automated decision making and assisted decision making by public sector bodies' (July 2021), The Legal Education Foundation, available at https://research.thelegaleducationfoundation.org/wp-content/uploads/2021/07/FINAL-Legal-and-Regulatory-Frameworks-Governing-the-use-of-Automated-Decision-Making-and-Assisted-Decision-Making-by-Public-Sector-Bodies-1.pdf.

automated", we mean that the output of the system is the sole basis of the decision and there is no "human in the loop", beyond the original design of the system or contributions to the system's inputs.

24. Decisions can also be partially automated. In one example of partial automation, the automated system may provide decision support in the form of additional information to aid a human decision-maker. An example of this might be a system that assesses whether an offender poses a risk of reoffending, generates a score, and presents that risk score to a parole officer to inform their decision. In a second example, the automated system may be streaming or triage system, which determines (or partially determines) the type or quality of human judgment required in a particular case. For example, a system could categorise a visa application as high risk, which means that the application is directed to a more senior official and subjected to a higher level of scrutiny. In our view, fully automated decisions should fall within the scope of Article 22. At first blush, it would seem that partially automated decisions should not. However, this is complicated by the problem of automation bias: a well-established psychological phenomenon whereby people put too much trust in computers.[12] This may mean that officials over-rely on automated decision support systems and fail to exercise meaningful review of an algorithm's outputs.

25. Article 22, properly defined, should prohibit *de facto* solely automated decision-making where, due to automation bias or for any other reason, the human official is merely rubber-stamping a score, rating or categorisation determined by the computer. It should require meaningful human oversight, rather than a token gesture. This may necessitate, for example, training to combat automation bias or other mitigations.

[12] See, for example, L.J. Skitka and others, 'Does automation bias decision-making?'(1999) 51 International Journal of Human-Computer Studies 991. For an example of automation bias in action in the UK, see the Independent Chief Inspector of Borders and Immigration, 'An inspection of entry clearance processing operations in Croydon and Istanbul: November 2016 – March 2017' (July 2017) at 3.7, 7.10 and 7.11, available at An-inspection-of-entry-clearance-processing-operations-in-Croydon-and-Istanbul1.pdf (publishing.service.gov.uk).

26. We would strongly resist a narrow definition of "legal" or "similarly significant" effects. In our own work, we have come across a number of decisions that we be understand to be encompassed by this term. This includes decisions about:

    - Qualifications. For example, the allocation of GCSE and A-level grades using Ofqual's standardisation model, intended for use in Summer 2020. Even if this model would not have directly affect students' legal rights, the effect on their long and short term prospects would have been sufficiently significant to fall within the scope of Article 22.[13]

    - Immigration enforcement investigations. For example, decisions about whether to investigate a couple who has given notice to be married to identify or rule out sham activity, made using the Home Office's 'triage tool'.[14] To be clear, we consider that not only the decision made *at the end of an immigration enforcement investigation*, but also *the decision to investigate in the first place* has legal or similarly significant effect for the purposes of Article 22.

    - Benefits sanctions investigations. For example, decisions to conduct benefits fraud or error investigations, made using the Department of Work and Pensions' "automated risk analysis and intelligence system".[15]

27. In our view, these sorts of decisions should remain within the scope of Article 22, and the fact that they fall within its scope should arguably be made clearer. For example, it may be useful to have a list within the UK GDPR indicating the sorts of decisions that have

---

[13] For further analysis of Ofqual's model, see Jack Maxwell and Joe Tomlinson, 'Model students: why Ofqual has a legal duty to disclose the details of its model for calculating GCSE and A level grades' (28 July 2020), available at https://ukconstitutionallaw.org/2020/07/28/jack-maxwell-and-joe-tomlinson-model-students-why-ofqual-has-a-legal-duty-to-disclose-the-details-of-its-model-for-calculating-gcse-and-a-level-grades/.

[14] For further analysis of the Home Office's triage tool, see Tatiana Kazim, 'Home Office refuses to explain secret sham marriage algorithm' (21 July 2021), available at https://www.freemovement.org.uk/home-office-refuses-to-disclose-inner-workings-of-sham-marriage-algorithm/.

[15] PLP's freedom of information request on this and the response is available here https://www.whatdotheyknow.com/request/automated_risk_analysis_and_inte#incoming-1661679.

legal or similarly significant effects. However, if Article 22 is to be adequately future-proofed and capable of accommodating unforeseen applications of new technologies, any such list must be non-exhaustive.

28. The consultation asks: *Q1.5.17. To what extent do you agree with the Taskforce on Innovation, Growth and Regulatory Reform's recommendation that Article 22 of UK GDPR should be removed and solely automated decision making permitted where it meets a lawful ground in Article 6(1) (and Article 9-10 (as supplemented by Schedule 1 to the Data Protection Act 2018) where relevant) and subject to compliance with the rest of the data protection legislation?*

29. PLP strongly disagrees with the proposal to remove Article 22 of the UK GDPR. See our comments above, in relation to Q1.5.14.

Chapter two – Reducing burdens on businesses and delivering better outcomes for people

Data Protection Impact Assessments

30. The requirement to undertake a Data Protection Impact Assessment (DPIA) is provided for under section 64 of the DPA 2018. Section 64(1) requires that a DPIA must be carried out "[w]here a type of processing is likely to result in a high risk to the rights and freedoms of individuals". Under section 64(3), the DPIA must include: "(a) a general description of the envisaged processing operations; (b) an assessment of the risks to the rights and freedoms of data subjects; (c) the measures envisaged to address those risks; (d) safeguards, security measures and mechanisms to ensure the protection of personal data… taking into account the rights and legitimate interests of the data subjects and other persons concerned." This requirement applies to all data controllers.

31. The consultation asks: *Q2.2.7. To what extent do you agree with the following statement: 'Under the current legislation, data protection impact assessment requirements are helpful in the identification and minimisation of data protection risks to a project'?*

32. PLP considers that DPIAs are an important tool for guarding against some of the risks posed by ADM systems, including discrimination. They can help with the identification and minimisation of such risks before they arise. This benefits providers and users of ADM technologies in that it helps them to avoid wasting resources on developing systems that flout legal standards protecting individual rights. It benefits individuals in that it helps to prevent them from being subject to rights-violating systems.

33. In our experience, DPIAs appear to have been useful in helping government departments to think through the full implications of using a given ADM technology and to help avoid legal challenges.

34. For example, the DPIA was valuable in understanding the implications of an ADM system piloted by the Ministry of Justice (MoJ) for use in the prison categorisation of newly sentenced offenders, as part of its Digital Categorisation Service (DCS).

35. Documents obtained by PLP under the Freedom of Information Act 2000 (FOIA) indicate that the ADM system works follows: the DCS contains a list of newly sentenced offenders in need of initial categorisation, drawn from the Prison National Offender Management Information System (P-NOMIS). The prison officer selects a prisoner, and the DCS automatically populates an online form with information about the prisoner, also drawn from the P-NOMIS. The officer can add further information as required. The algorithm uses all this information to generate a 'provisional category' for the prisoner, which the officer can accept or reject.

36. The MoJ identified a risk that the DCS may be over-categorising BAME prisoners: during the DCS pilot, BAME prisoners were initially categorised as Category B at a higher rate than white British prisoners: 10% as compared to 7%.

37. Most of this information was set out in the DPIA. Generally speaking, disclosure of the DPIA can help individuals subject to an ADM system, as well as organisations that support them, to understand the operational details and risk of discrimination posed by that system.

This, in turn, can better equip individuals and organisations to hold the users of ADM systems to account.

38. Yet, as Swee Leng Harris has noted, there are no legal requirements for public disclosure of DPIAs[16] – beyond the general requirement in the FOIA to disclose information in response to a request, provided that no exemption applies. In PLP's experience, requests for disclosure of DPIAs are not always granted. Hence, while we consider that DPIAs are a useful protection, our view is that more could be done to strengthen their utility.

39. We agree with Swee Leng Harris that, "transparency and accountability would be enhanced if DPIAs by government were published by default, so that there was a publicly accessible database of government DPIAs."[17] Mandatory publishing of government DPIAs on a centralised, searchable and publicly accessible database would help to uphold rule of law standards. This could also tie in with the compulsory transparency reporting on the use of algorithms in public decision-making proposed in Chapter 4 of the consultation: a link to the relevant DPIA could be one of the pieces of information made available (see below).

40. Further, while it is valuable to have a forward-looking DPIA in order for risks to be identified and mitigated early on, this is also a limitation. A forward-looking DPIA cannot account for risks that only become clear once the system is deployed. Under Article 35(11) GDPR, a review should be undertaken to check whether data processing is being performed in accordance with the DPIA "[w]here necessary" and "at least when there is a change of the risk represented by processing operations". However, this seems to put the cart before the horse: change of risk may be difficult to identify without carrying out a review. In our view, a requirement for routine DPIAs, or a similar form of evaluation, throughout the period of

[16] Swee Leng Harris, 'Data Protection Impact Assessments as rule of law governance mechanisms' (30 March 2020), Cambridge University Press, available at https://www.cambridge.org/core/journals/data-and-policy/article/data-protection-impact-assessments-as-rule-of-law-governance-mechanisms/3968B2FBFE796AA4DB0F886D0DBC165D.

[17] Ibid.

deployment of an ADM system ought to be considered. A measure of this kind may help to identify unanticipated rights-violations.

41. The consultation asks: *Q2.2.8. To what extent do you agree with the proposal to remove the requirement for organisations to undertake data protection impact assessments?*

42. PLP disagrees with the proposal to remove the requirement for organisations to undertake DPIAs. As explained above, we consider that DPIAs have value both as forward-looking tools, helping organisations to identify and minimise risks to the rights and freedoms of data subjects before they arise, and – where DPIAs are disclosed – as tools for promoting transparency and accountability. We consider that more should be done to enhance the utility of DPIAs. They should not be done away with.

43. As for the concern that the requirement to undertake a DPIA is too burdensome, a controller is only required to carry out a DPIA where "a type of processing is likely to result in a high risk to the rights and freedoms of individuals".[18] This is a high bar, and one that the majority of data processing is unlikely to meet. To the extent that this requirement gives all data controllers pause for thought – even those who ultimately decide that they are not required to undertake a DPIA – this is a positive. Data use, and the use of new technologies more generally, should not be presumed to be beneficial. Rules like the requirement to undertake a DPIA are important breaks on the process.

44. That the requirement to undertake a DPIA is pegged to the level of threat posed to the rights and freedoms of individuals, rather than to the circumstances of the data processer, is in our view entirely appropriate. The single regulatory standard provides a minimum level of protection for all individuals.

[18] Article 35(1) of the UK GDPR.

Data subject access requests

45. Data subject access requests are provided for under Article 15 of the UK GDPR. They allow individuals to check the accuracy of their personal data, learn more about how their data is being used and with whom their data is being shared, and obtain a copy of the data held about them. They are also an important investigative tool, helping to ensure transparency and uncover and prevent maleficence.

46. The consultation proposes two key changes based on provisions under the FOIA: first, allowing subject access requests to be refused if the cost of answering the request exceeds particular limits (under the FOIA, this is currently £450 – equivalent to 18 hours at £25 per hours – for local authorities and £600 – equivalent to 24 hours at £25 per hour – for central government); and, second, permitting burdensome requests to be refused as vexatious. The consultation also proposes to impose a "small nominal fee" on data subjects, a requirement which is no part of the FOIA.

47. The consultation asks: *Q2.3.3. To what extent do you agree that introducing a cost limit and amending the threshold for response, akin to the Freedom of Information regime (detailed in the section on subject access requests), would help to alleviate potential costs (time and resource) in responding to these requests?*

48. PLP considers that reduction in the cost of responding to data subject access requests is not a sufficient justification for curtailing the rights of data subjects. We strongly disagree with the proposals to introduce a cost limit and amend the threshold for response.

49. While FOI requesters can complain to the ICO if they are refused on a costs basis, the ICO does not generally enforce data subject access rights.[19] Therefore, a data subject would have limited avenues for redress if the cost limit or threshold is incorrectly applied.

[19] See Campaign for Freedom of Information, 'Rights to see personal data at risk' (10 September 2021), available at https://www.cfoi.org.uk/2021/09/rights-to-see-personal-data-at-risk/.

50. Further, and perhaps more importantly, the position of a data subject is profoundly different from the position of a requester under the FOIA. The consultation document acknowledges that the ability of an individual to access their own data is a fundamental right. In our view, measures constituting an acceptable compromise under the FOIA would not, therefore, be acceptable in relation to the rights of data subjects.

51. The consultation asks: *Q2.3.4. To what extent do you agree with the following statement: 'There is a case for re-introducing a small nominal fee for processing subject access requests (akin to the approach in the Data Protection Act 1998)'?*

52. PLP considers that the case against introducing a fee for processing data subject access requests is stronger than the case for it. We strongly disagree with the proposal to introduce a fee.

53. This is not least because the proposal appears to disregard the unique position of data subjects (explained above) and, perversely, would put them at a disadvantage compared with requesters under the FOIA – though, for the reasons set out below, PLP does not consider that this discrepancy could be resolved by 'levelling down', imposing a fee on both data subjects and FOIA requesters.

54. For someone struggling to make ends meet, a 'small nominal fee' may be unaffordable. In 2019/20, 22% of all people and 25% of working-age people with children were living in poverty.[20] For these people, a fee is likely to be prohibitive, resulting in a serious chilling effect on requests.

55. Further, studies have shown that some protected characteristics including race, sex (in the case of single mothers), and disability are associated with an increased risk of poverty.[21]

---

[20] Joseph Rowntree Foundation, 'Poverty rate by person type over time, after housing costs', available at https://www.jrf.org.uk/data/poverty-rate-person-type-over-time-after-housing-costs-ahc.

[21] See, for example, Sara Davies and David Collings, 'The Inequality of Poverty: Exploring the link between the poverty premium and protected characteristics' (February 2021), University of Bristol Personal Finance Research Centre, available at https://fairbydesign.com/wp-content/uploads/2021/02/The-Inequality-of-Poverty-Full-Report.pdf.

The more protected characteristics someone holds, the greater their statistical risk of poverty.[22] This means that the chilling effect of the fee would very likely be more severe in respect of requests by people from disadvantaged and marginalised groups. This would be especially egregious given that some known ADM systems, such as immigration enforcement systems and welfare fraud and error detection systems, will likely have a disproportionate effect on these groups. To give an example, the proposal would mean that while a disabled person may be more likely to be processed by a welfare fraud detection system, they will be less likely to be able to afford the fee to find out about and challenge this.

56. In short, the proposal to introduce a fee – especially in conjunction with the proposals to introduce a cost limit and amend the threshold for response – would curtail the rights of data subjects in a way that is unfair and unjustified.

57. It should be noted that the compulsory transparency reporting envisaged in Chapter 4 would not be a complete solution to this problem. While this measure would provide general information about ADM systems, it would not provide specific information about individual data subjects and the way their personal data has been processed. We disagree with the proposals relating to data subject access requests regardless of whether compulsory transparency reporting is introduced.

Chapter four – Delivering better public services

Government algorithms and compulsory transparency reporting

58. PLP has found that a lack of transparency in government use of algorithms – or, to use our preferred terminology, ADM systems – is a persistent problem.

[22] Ibid.

59. Opacity in the use of ADM systems may be intentional, negligent, or due to the complexity of the system. The latter is a particular problem when it comes to machine learning. A machine learning algorithm may be a 'black box', even to an expert.

60. Intentional opacity can occur where government, or a private developer contracted by government, deliberately withholds information about the system due to concerns about commercial confidentiality or possible abuse and circumvention of the system's rules. A number of our FOI requests have been refused on the basis of section 31 FOIA: releasing the information would prejudice the ability to detect fraud and crime.

61. Even where our FOI requests have been granted and substantial amounts of information disclosed, relying on such requests as means of identifying government use of ADM systems is deeply unsatisfactory and the successful requests go little way to mitigating the overall opacity. First, the twenty day statutory time limit is often extended and delays in responses frequently last for several months. Second, it is often only at the internal review stage that meaningful information is disclosed, if ever. Both of these factors mean that relying on FOI requests is a slow and cumbersome method of finding out about ADM systems. Third, successful FOI requests usually require knowing where to look. There must usually be some kind of 'lead' in a publicly available document. Requests framed in more general terms, such as a request for a list of Home Office systems using a streaming tool to separate applications into different categories[23], or a request for strategy documents and DPIAs in relation to the use of algorithmic decision-making by a government department,[24] are likely to be refused on the basis of the cost limit under section 12 FOIA.

[23] See PLP's request, available at https://www.whatdotheyknow.com/request/streaming_tool#incoming-1634835.

[24] See PLP's request, available at https://www.whatdotheyknow.com/request/algorithmic_decision_making#incoming-1662209.

62. Opacity is arguably a cost in and of itself, but also comes with costs in terms of the ability of people to hold the state to account. We consider that transparency is a prerequisite for accountability (see further below, in response to Q4.4.1.).

63. Moreover, the current opacity around the use of ADM systems may well contravene administrative law doctrines including the principle of transparency, procedural fairness, the duty to give reasons, and the duty of candour.[25] Each of these four principles apply at different stages of the decision-making process but all operate to make exercises of public power more transparent. For example, procedural fairness governs the stage where the government proposes to exercise a power in a particular case. Generally speaking, it requires that a person whose rights and interests will be affected by the exercise of power be given the gist of the case against them, so that they can make meaningful representations before the decision is made. In cases where an ADM system is used, this may require disclosure of a fully executable version of the system.[26]

64. The consultation asks: *Q4.4.1. To what extent do you agree that compulsory transparency reporting on the use of algorithms in decision-making for public authorities, government departments and government contractors using public data will improve public trust in government use of data?*

65. In principle, PLP supports compulsory transparency reporting. This is subject to two important caveats: first, this must be meaningful transparency; and, second, transparency alone is not enough. Subject to these two caveats, we consider that compulsory transparency reporting has the potential to help government earn public trust in their use of data.

66. The first caveat is addressed below, in our response to Q4.4.2. As for the second caveat, transparency is a good start, but it is only one step along the road to trustworthy use of

[25] Jack Maxwell and Joe Tomlinson, 'Government Models and the Presumption of Disclosure' (July 15, 2020), available at https://ssrn.com/abstract=3652602.

[26] *R (Eisai Ltd) v National Institute for Health and Clinical Excellence* [2008] EWCA Civ 438.

ADM technology. Amongst other things, accountability is also essential. To an extent, accountability is dependent on transparency.[27] However, it goes beyond transparency, in that it requires adequate avenues for people to challenge the development and deployment of ADM systems, together with effective enforcement mechanisms and the possibility of sanctions. Definitions of accountability differ. But it has been suggested that any adequate definition will involve three elements:

i. *Responsibility* for actions and choices. There must be an accountable party who can be praised, blamed, and sanctioned.

ii. *Answerability*, which includes: first, capacity and willingness to reveal the reasons behind decisions to a selected counterpart (this could be the community as a whole); and, second, entitlement on the part of the counterpart to request that the reasons are revealed.

iii. *Sanctionability* of the accountable party, where 'sanctions' range from social opprobrium to legal remedies.[28]

67. PLP considers that this is a useful starting point in thinking about how to achieve accountability in an ADM context.

68. As well as transparency and accountability, we also endorse the following principles: anti-discrimination; reflexivity (involving continuously review of the way people's beliefs and judgments may influence the development and use of the system); and respect for privacy and data rights.[29] In order to put these principles into practice and promote trustworthy ADM technologies, compulsory transparency reporting should not be considered in

---

[27] Michele Loi and Matthias Spielkamp,'Towards accountability in the use of Artificial Intelligence for Public Administrations' (21 July 2021), available at https://algorithmwatch.org/en/wp-content/uploads/2021/05/Accountability-in-the-use-of-AI-for-Public-Administrations-AlgorithmWatch-2021.pdf.

[28] Ibid.

[29] See also our response to the Justice and Home Affairs Committee's call of evidence on the use of new technologies and the application of the law, available at https://committees.parliament.uk/writtenevidence/39761/pdf/.

isolation from other proposals in the consultation document. Instead, compulsory transparency reporting should be implemented alongside other suggestions we have made throughout this response: the DPA 2018 and UK GDPR should be amended to make clear that data processing which leads to breaches of the EA 2010 is unlawful; Article 22 should be retained and its key terms clarified to ensure its broad practical application; the requirement to undertake DPIAs should be retained and strengthened; and the rights of data subjects to make requests for information should remain unattenuated.

69. The consultation asks: *Q4.4.2. Please share your views on the key contents of mandatory transparency reporting.*

70. The EU Commission's proposed artificial intelligence (AI) regulation, adopted on 21 April 2021,[30] would include a public register of high risk AI systems in the form of a database, managed by the EU Commission, to which AI providers would be obliged to provide meaningful information about their systems. This proposal points towards three key areas for consideration when it comes to the contents of compulsory transparency reporting in the UK: first, the types of system to be included; second, the duty-bearers; and third, the type of information to be disclosed. Of course, analysis of the EU Commission's proposed AI regulation of this is still unfolding, and it certainly should not be taken as the gold standard. However, we consider it to be a useful comparator.

71. Regarding the first consideration, the EU Commission's proposed AI regulation uses 'high risk' as the touchstone. This is fleshed out by way of a list, in Annex III, of high risk systems organised into eight (unamendable) categories. These categories include education and vocational training, employment, access to essential public and private services and benefits, and law enforcement. Regarding the second consideration, duty-bearers include not only public bodies and government contractors, but also private parties; indeed, any 'providers' of AI systems. As for the third consideration, the information that would be

[30] The proposal in available at EUR-Lex - 52021PC0206 - EN - EUR-Lex (europa.eu).

available in the EU database is set out in Annex VIII and includes: details of the provider; the trade name and other identifying details of the system; a description of the intended purpose of the system; its status (for example, whether it is currently in use); and electronic instructions for use.

72. We appreciate that the volume of ADM systems in use by public bodies in the UK – as well as by private parties – is likely to be very high. A significant proportion of these will likely perform routine administrative tasks. Therefore, we recognise that it may not be helpful or proportionate for all of these systems to be included in a public register or database. It may be that a touchstone like 'high risk' is more appropriate.

73. We note that Karen Yeung et al have argued that the definition of 'high risk' in the context of the EU Commission's proposed AI regulation requires improvement.[31] For example, they suggest that emotion recognition systems should have been included in the list of high risk systems. Further, they suggest that the categories in Annex III – which limit the types of system that can be considered high risk – should be subject to amendment, to ensure the legislation is adequately future-proofed.

74. With this in mind, PLP suggests that compulsory transparency reporting requirements in the UK should apply at minimum to high risk systems, broadly and flexibly defined. The possibility of including a non-exhaustive list of high risk systems could be considered.

75. At present, DCMS's proposal for compulsory transparency reporting is limited to public authorities, government departments and government contractors using public data. Given that, by definition, high risk systems pose a threat to individual rights no matter who they are provided or used by, PLP suggests that DCMS should consider whether other parties

[31] Nathalie A Smuha, Emma Ahmed-Rengers, Adam Harkens, Wenlong Li, James MacLaren, Riccardo Piselli, and Karen Yeung, 'How the EU Can Achieve Legally Trustworthy AI: A Response to the European Commission's Proposal for an Artificial Intelligence Act' (August 5, 2021), available at SSRN: https://ssrn.com/abstract=3899991.

providing or using high risk ADM systems should be subject to compulsory transparency reporting requirements, too.

76. In order to achieve meaningful transparency, PLP suggests that the following types of information should be considered for inclusion in compulsory transparency reporting requirements: details of the provider and user(s) of the ADM system; identifying details of the ADM system (such as its name); purposes for which it is used; status, including period and scale of deployment; any training data, methodologies, and techniques; DPIAs; Equality Impact Assessments (EIAs); an explanation of how the system works; and an executable version of the system.

77. The consultation asks: *Q4.4.3. In what, if any, circumstances should exemptions apply to the compulsory transparency reporting requirement on the use of algorithms in decision-making for public authorities, government departments and government contractors using public data?*

78. PLP recognises that some exemptions may be necessary. It may be that the types of information and level of detail that can be provided without jeopardising the public interest will vary according to the area in which the ADM system is used. For example, it may be that less detail can be provided about law enforcement systems. An exemption along these lines would chime with Annex VIII of the EU Commission's proposed AI regulation, under which electronic instructions for use are not to be included in the database where the system is used for law enforcement or migration, asylum and border control management.

79. An obvious model for any such exemptions would be the FOIA. However, PLP has some reservations about this approach. We have found that there can be over-reliance on the exemption under section 31 FOIA, for example. Public authorities can be too quick to decide that the public interest falls in favour of withholding information.

80. At the very least, we do not think that exemptions should apply to the disclosure of high-level information, including the fact that there is an ADM system in use in a particular

context. If more detailed information about a particular system is withheld on the basis of an exemption, there should be readily accessible avenues for challenging this, with the possibility of review by an independent regulator such as the ICO.

19 November 2021