Withers LLP response to UK Government's White Paper "*A pro-innovation approach to AI regulation*"

1. Introduction

1.1 Withers LLP is an international law firm which advises a range of clients, including businesses, private individuals and organisations, all with a burgeoning interest in AI and how it is shaping and transforming the world around them.

1.2 We have an established international dispute resolution practice acting in complex multi-jurisdictional cases worldwide. We are well known in the art, sport, fashion, luxury assets and natural resources sectors, with experience across various specialisms including international arbitration, public international law, tax controversy, private wealth disputes, divorce and family, media and reputation management, employment and regulatory, white collar investigations, professional negligence, board and shareholder disputes, insolvency and fraud. Representing businesses, partnerships, individuals and families, we strive to achieve outstanding results in any dispute – be that litigation, arbitration, negotiation, trials or appeals. Our team regularly provides corporate, dispute resolution, and regulatory and compliance services in the cryptocurrency and fintech space. We are experienced with the existing and changing regulatory framework applicable to individuals and entities operating in or adjacent to this growing area, and we represent individuals and entities, by way of example, in cryptocurrency and digital asset related matters across the globe.

1.3 We also advise entrepreneurs, disruptive start-ups and early stage clients in the tech sector through our Withers Tech team, which works with early stage businesses, entrepreneurs and venture capital investors. Many of these clients are already developing or investing in AI products and systems as part of their core businesses or new products in development.

1.4 This synergy of interests across our client base gives us a position of strength to advise clients on the relative risks and merits of the deployment of AI technology in relation to their business and/or personal arrangements. Similarly, this experience also places the firm in a strong position to give a meaningful and constructive response to the UK Government's White Paper on AI, *A pro-innovation approach to AI regulation* (the 'White Paper').

1.5 This response focuses on the White Paper principles (each a 'Principle') and the specific legal queries posed by the White Paper, where our expertise and experiences enable us to provide a constructive response. We also note that the White Paper is, by design, a high-level document and that the necessary levels of granularity on certain concepts have not yet been developed into prospective policy, guidelines etc. In relation to this, where applicable, we have looked to make some observations on how the UK Government may wish to approach adding detail to such concepts in due course.

2. High-Level Summary

2.1 AI is an incredibly dynamic and rapidly evolving field. The UK Government at this stage has not outlined an overly rigid regulatory structure which risks becoming rapidly outdated or obsolete, as AI technology and uses of technology develop. Consequently, the UK Government still needs to outline in detail how its proposed regulatory framework will work in practice and the standards that businesses using and developing AI will need to meet.

2.2 By way of example, we have already illustrated to clients of the firm (such as employers, recruiters and insurers) salient points arising out of the White Paper and how it might impact their business practices. In particular, we have emphasised risks arising out of different approaches to privacy across jurisdictions and how those may translate in potential regulatory issues in relation to using AI tools in (i) e-recruitment practices and (ii) assessing insurance premiums.

2.3 From the perspective of early stage businesses, there is a question of the extent to which compliance with a regulatory framework will be mandatory, on the basis that businesses at this stage may struggle to comply due to budgetary and expertise constraints. Conversely, there is a risk that mandatory compliance may limit these early stage businesses from entering the AI market on the basis that the steps for compliance are too onerous or require inaccessible expertise.

2.4 It is unclear what the legal consequences will be for businesses using AI technology during the regulatory framework's initial "non-statutory" phase.

2.5 We note that as part of its approach, the UK Government intends to produce a regulatory roadmap to set out deadlines by which different regulators will need to publish their respective guidance. This timeline will need to clarify what standards need to be met, when, and by whom.

2.6 If the UK Government delays introducing regulation, there is a risk that it will be too difficult to reverse AI technologies that are embedded in industries, to the extent such technologies are not fully compliant with such regulation.

2.7 Noting the UK Government's "pro-innovation" approach, we note that a highly regulated AI sector risks giving too much control to a handful of large tech companies, which may be the only companies with the financial resources to operate in a highly regulated space. It will be critical to ensure the proportionality and/or reasonableness of any regulatory overlay onto a fledging new business.

2.8 We note that the FCA in the past has deployed a regulatory sandbox approach to fintech businesses which has aided both regulators' and new businesses' understanding of what the appropriate regulatory overlay is to such new and innovative businesses. We understand that the Government envisages running similar sandbox schemes which would close the feedback loop between regulators and businesses and aid the UK Government in implementing its desired iterative and pro-innovation approach.

2.9 The UK Government may also wish to impose restrictions on what the most powerful versions of AI technologies can be used for. For example, certain uses may be restricted for use only in a national security or policing context and will not be available to everyday businesses or consumers.

3. Responses to White Paper consultation questions

3.1 Our responses to the White Paper consultation questions are outlined below. For ease of reference, the numbering reflects the numbering used for the questions in Annex C of the White Paper.

The revised cross-sectoral AI principles

1. *Do you agree that requiring organisations to make it clear when they are using AI would adequately ensure transparency?*

1.1 Yes, provided that organisations are also required to communicate: (i) how, and (ii) for which purposes they are using AI in a standardised way so that – in the case of customers or clients – they have the opportunity to make an informed decision as to whether or not they consent to this use of AI. As to the standard form of communication: the UK Government should provide minimum requirements necessary for compliance.

1.2 Given that the use of AI is so ubiquitous, smaller or less sophisticated organisations should not be exempt from such a requirement *a priori*. However, in devising guidelines and transparency standards, the UK Government should consider introducing additional/higher transparency requirements for those organisations that deploy high-risk AI tools.

1.3 The UK Government should also consider what the outcome of this requirement would look like. We anticipate that this would likely be a legal notice on an organisation's website, which risks users of AI-powered tools and services either (i) ignoring entirely, or (ii) becoming desensitised to (in the same way that cookie notices are now ubiquitous and largely ignored by web users). If a legal notice

on a website is to be applied, the level of compliance with minimum standard formatting requirements (unlike cookie setting banners that come in all shapes and sizes and with differing levels of granularity) should also be considered. If this approach was taken, unlike website cookie notices, there is an opportunity to facilitate these notices at the genesis of the technology in a consumer setting, which may increase the prospect of it being developed in a user-friendly and practical way, as opposed to notices being retrospectively required as part of regulatory compliance requirements.

2. *What other transparency measures would be appropriate, if any?*

2.1 The UK Government should consider whether organisations should be required to undertake an appropriate AI risk assessment (potentially similar to those developed by the ICO) and provide an explanation of measures they have taken to mitigate the relevant risks.

2.2 There is an argument that risk assessment reports should be provided in a standard form (i.e. they should comply with minimum requirements applicable to all organisations) so that users of AI, auditors and the like can compare different organisations easily and ensure that all organisations are considering the same categories of risk (with of course the option of not expanding on some where these are not applicable to the organisation in question).

2.3 Regulators should work together in order to produce said form – this will also ensure greater cross-collaboration between them. In particular, it will be important for any risk assessment report to show continuous assessment by the organisations – otherwise there is a risk that they will be reverse-engineered to explain ex-post what the organisation's process is.

2.4 We note that the UK Government has launched an "Assurance Portfolio" toolkit and that international bodies like the OECD are compiling a catalogue of tools and metrics for trustworthy AI. It will therefore be important – for the burgeoning AI assurance industry, which the UK wants to capitalise on – for risk assessments to be properly compared by auditors.

2.5 If assessments are not in a standard form, then there is a real risk that companies will: (i) not be equipped to consider issues that they should have taken into account; or (ii) deliberately avoid considering certain factors in an attempt to minimise potential harm. The risk assessment form should however allow companies to include any additional issues in narrative form (as the ICO toolkit envisages) so as to not limit their ability to consider wider points.

2.6 If there is a significant risk of harm to users affected by AI tools and services as a result of an organisation's use of AI (e.g. in health-related sectors), then organisations should be required to make AI users aware of this fact in advance within their communication notice and include their compliance with sector-specific guidance (e.g. additional MHRA guidance in relation to health products).

2.7 Organisations should also be required to provide information to AI users on how they may contest and/or obtain redress for harms which are caused by the use of AI.

3. *Do you agree that current routes to contestability or redress for AI-related harms are adequate?*

3.1 No. Current routes to contestability or redress for AI-related harms are not adequate as – until regulators clarify existing routes (including informal ones) as envisaged in the White Paper – these remain primarily grounded in the UK GDPR. Whilst the UK GDPR does offer a route to contestability or redress for automated decision making, this has not proven sufficiently clear or effective.

3.2 By way of example, what counts as sufficient human intervention (so that a decision cannot be said to be "solely automated") in circumstances where a human actor may be greatly influenced by an AI's outcome has been the subject of debate and litigation.

3.3 Further, many companies use dark patterns and other manipulative strategies in order to obtain an individual's explicit consent to automated decision-making. Even with an individual's explicit consent, companies are required to introduce simple ways for individuals to request human intervention/challenge a decision.

3.4 It is also difficult for individuals to assess whether an organisation is deploying AI tools which systematically exhibit, for example, bias and discrimination as individuals due to information asymmetry.

4. *How could routes to contestability or redress for AI-related harms be improved, if at all?*

4.1 Consumers and users affected by AI tools and services should be able to complain – as also envisaged by the CMA – to authorities like regulators with investigatory powers in addition to bringing regular actions before the English courts. Having additional routes of contestability is critical because of the high cost associated with bringing court claims which could deter individuals from presenting such an action.

4.2 The UK Government should consider whether or not a new ombudsman or existing ombudsmen could take on the task of creating easier routes to report and contest AI-related harms whilst also seeking to avoid a situation in which floods of spurious claims are brought (for example, a wave of automated claims being made by bots).

4.3 All relevant regulators should issue clear guidance on what actions, or types of action, by a business using AI can be contested in light of the principles.

4.4 Clearer guidance should be given (e.g. by the courts and relevant regulators) as to what constitutes automated decision-making or sufficient human intervention by providing more tailored examples falling within scope of relevant regulators; these should then be expanded where necessary following the sandbox programme (whose aim is to reduce the feedback loop between industry and UK Government).

5. *Do you agree that, when implemented effectively, the revised cross-sectoral principles will cover the risks posed by AI technologies?*

5.1 Presently there is a wider lack of clarity on: (i) what standard is required by the existing principles, and (ii) what would constitute 'effective implementation' of the existing principles which makes it difficult to assess whether further cross-sectoral principles are needed.

5.2 The UK Government may wish to consider whether more cross-reference to environmental, social, and corporate governance ("ESG") reporting requirements should be made within the AI principles or vice-versa. AI technologies have the potential to harm to our democracy and environment and could run counter to sustainability goals and priorities (to be interpreted not just in relation to the environment but also, by way of example, in relation to the labour market).

5.3 The UK Government should also consider whether, in producing risk assessment reports, organisations using and deploying AI should also reflect on whether or not their use of AI is necessary or beneficial by reference to the risk profile of the relevant AI tool and also wider ramifications relating to, for example, the labour market.

6. *What, if anything, is missing from the revised principles?*

6.1 It is unclear how the principles interact with each other. More specifically, it is not clear how the proposed framework would deal with a situation where either: (i) two or more principles impose conflicting responsibilities, and/or (ii) conflicting responsibilities are imposed by one or more principles, and wider law/regulation.

6.2 For example: (i) the principles of explainability and fairness may come into conflict with one another as what is considered "fair" may vary greatly from group to group and case to case and it may be

possible for a "fairer AI system" to be less explainable than one that is considered to be less "fair"; and (ii) the principle of transparency and explainability may come into conflict with laws and regulations concerning the right to privacy.

6.3 We further note that the White Paper envisages that when carrying out risk assessments, a company may need to consider the risk of not using AI in a particular scenario. This highlights the types of conflicts that may arise in the future if further clarity is not provided by the regulators.

A statutory duty to regard

7. *Do you agree that introducing a statutory duty on regulators to have due regard to the principles would clarify and strengthen regulators' mandates to implement our principles, while retaining a flexible approach to implementation?*

7.1 We do not think that introducing a statutory duty on regulators to have due regard to the principles would clarify regulators' mandates. This is because: (i) it seems likely that there would be an overlap in the mandates of the various regulators, and (ii) a statutory duty does not, in of itself, add clarity to what the principles require, nor does it provide the regulators with the expertise required.

7.2 However, a statutory duty could possibly strengthen regulators' mandates as it would require regulators to consciously consider the principles and ensure they have done enough to comply with their duty to do the same. This in turn would embolden regulators to take a pro-active approach in implementing the principles and to issue further guidance where appropriate in keep with the flexible approach to implementation.

8. *Is there an alternative statutory intervention that would be more effective?*

8.1 An alternative approach would be a fully centralised "new" AI regulator, provided with: (i) a statutory duty to have due regard to the principles, and (ii) the necessary in-house staff expertise, and working with the current regulators.

8.2 Alternatively, the UK Government could take additional steps to ensure that each regulator has a dedicated AI team whose members will be required to collaborate and meet on a regular basis with their respective counterparts in other regulating bodies. The dedicated AI teams would also routinely meet with industry, academia and civil society to reduce the feedback loop given that the UK's proposed approach is intended to be flexible and iterative. This would create a hybrid model between the UK Government's current proposed approach and the new standalone AI-regulator model.

New central functions to support the framework

9. *Do you agree that the functions outlined in section 3.3.1 would benefit our AI regulation framework if delivered centrally?*

9.1 The proposed new central functions outlined in section 3.3.1 of the White Paper may not necessarily have the desired effect of the UK Government - namely to ensure a light-touch approach to promote innovation. In fact, these central functions may even prove to be counter-productive, as a result of friction between the UK Government's proposed new central coordinating function and the individual regulators' existing individual functions.

9.2 This friction however may be greatly reduced if a new independent AI regulator and/or audit office is created and tasked with carrying out those central functions as, without a central driving force, momentum may be lost and the risk of contradictory guidance being issued by different regulators may increase. We note that creating a new independent AI regulator and/or audit office or empowering an existing regulator – like the ICO – to take such additional roles will require additional funding.

10. *What, if anything, is missing from the central functions?*

Deliberately blank.

11. *Do you know of any existing organisations who should deliver one or more of our proposed central functions?*

Deliberately blank.

12. *Are there additional activities that would help businesses confidently innovate and use AI technologies?*

Deliberately blank.

12.1 *If so, should these activities be delivered by government, regulators or a different organisation?*

Deliberately blank.

13. *Are there additional activities that would help individuals and consumers confidently use AI technologies?*

Greater education of individuals and consumers is required. Media and government officials unfortunately often use language and images that anthropomorphise AI and emphasise specific characteristics, often drawing from science fiction films and pop culture. This creates confusion as the AI tools that risk causing significant harm to individuals and consumers are invisible or embedded in applications we commonly use (e.g. facial recognition technologies). The UK Government should consider strategies to educate individuals and consumers not just on long-term low-probability high-risk situations (threat of mass extinction) but, more importantly, on short-term, high-probability risks and how to navigate these confidently.

For example, if a standardised communication notice is devised (to be applied on company websites etc if they use AI tools), individuals and consumers should be taught how to navigate the various options contained in the notice in practical ways. We note that Sweden has piloted an initiative called "AI Competence for Sweden" which, amongst other things, included launching free AI courses online. Similar courses could be replicated in the UK.

13.1 *If so, should these activities be delivered by government, regulators or a different organisation?*

These activities should be delivered by the UK Government or by an independent AI regulator/auditor. We would envisage it being essential to ensure that whoever is delivering the education modules has appropriate expertise.

14. *How can we avoid overlapping, duplicative or contradictory guidance on AI issued by different regulators?*

14.1 Different regulators should be required to communicate with each on a regular basis, and the boundaries of each regulator's remit should be clearly communicated to all regulators. See our suggested approach set out at *Question 8* above.

Monitoring and evaluation of the framework

15. *Do you agree with our overall approach to monitoring and evaluation?*

Deliberately blank.

16. *What is the best way to measure the impact of our framework?*

Deliberately blank.

17. *Do you agree that our approach strikes the right balance between supporting AI innovation; addressing known, prioritised risks; and future-proofing the AI regulation framework?*

Deliberately blank.

18. *Do you agree that regulators are best placed to apply the principles and government is best placed to provide oversight and deliver central functions?*

18.1 See response to *Question 9* above.

Regulator capabilities

19. *As a regulator, what support would you need in order to apply the principles in a proportionate and pro-innovation way?*

Deliberately blank.

20. *Do you agree that a pooled team of AI experts would be the most effective way to address capability gaps and help regulators apply the principles?*

Deliberately blank.

Tools for trustworthy AI

21. *Which non-regulatory tools for trustworthy AI would most help organisations to embed the AI regulation principles into existing business processes?*

21.1 We note that specific AI certifications / audits (which are referred to in the White Paper) may help organisations using AI to embed principles into existing business processes. We understand that, through the CDEI Portfolio of AI Assurance Techniques, the government is already working closely with industry members.

Final thoughts

22. *Do you have any other thoughts on our overall approach? Please include any missed opportunities, flaws, and gaps in our framework.*

22.1 Most businesses will be operating in both the UK and EU, so any business-friendly approach will facilitate organisations complying with both regimes in a single step, e.g. a single user-facing 'How we use AI' policy.

22.2 The UK Government's one-size fits all approach – unlike the EU's risk-based approach – does not account for the fact that only a relatively narrow portion of AI use cases will pose a significant risk to end users. For the majority of use cases (i.e. those which are designated by the EU framework as 'low/minimal risk'), the UK's approach will likely be less streamlined and more onerous than the approach taken in the EU. Put another way: (i) in the UK, the 'average' AI organisation may have a compliance burden disproportionate to the risk level associated with its activity, and (ii) the same business operating in the EU would potentially have a lower compliance burden.

22.3 Whilst there are riskier categories of AI which will require more stringent laws and regulation (for example, use of AI in the military), it is also important to recognise that what appears to be "typical" AI usage may – if not regulated properly – still have serious implications for end users. The UK Government's principle-based approach appropriately envisages this. However, in order to ensure greater inter-operability between jurisdictions, which will be key to international cooperation and trade, the UK Government needs to engage with key players in the international community.

22.4 If the UK Government intends to require (as part of its endeavours to build an AI assurance service industry and to build public trust in AI tools and services) that organisations using and deploying AI tools (e.g. in financial services) be registered or authorised by the regulators then there may be an

issue of what to do with existing businesses that may at a future date have to be registered and/or authorised by regulators.

22.5 The general approach is to introduce an interim regime where existing businesses can carry on whilst applying for the required registration and/or regulatory authorisation with a future cut-off point by which businesses must be registered and/or authorised and thereafter they cannot carry on what will become regulated businesses without these registrations and/or authorisations.

22.6 However, based on our experience advising clients in the crypto asset sector, the above general approach requiring crypto asset businesses to register with the FCA had stark consequences for existing crypto asset businesses in the UK last year where some 200 applied for the FCA registration required by AML regulations and only some 40 were registered by the FCA i.e. some 20% by the 31 March 2022 cut off point with the remainder c80% forced offshore and/or to have to cease business in the UK.

22.7 The UK Government should therefore consider whether the general approach to introduce an interim regime is viable or whether a different approach would be more appropriate as regulation begins to apply to businesses using and deploying AI tools and services.

Legal responsibility for AI

23. *L1. What challenges might arise when regulators apply the principles across different AI applications and systems? How could we address these challenges through our proposed AI regulatory framework?*

23.1 Because the proposed principles are deliberately non-prescriptive, it is possible that they will be applied differently by different individuals – even within the same regulator. This could result in haphazard/inconsistent decisions being taken by the regulators, which: (i) would be unfair, and (ii) could unduly influence what business models are viable and which are not. These risks are exacerbated by the current situation, where there is a very large range of AI applications and systems. This challenge could potentially be overcome by upskilling staff within the relevant regulators.

23.2 Typically, commercial contracts will require the parties to comply with 'applicable laws'. In circumstances where the proposed AI regulation is not placed on a statutory footing, compliance with the proposed principles may be discretionary and so lead to inconsistent levels of compliance. This challenge could potentially be overcome by giving the principles a statutory footing.

23.3 Relying on the principles becoming industry-standard takes time and success is not guaranteed. The UK Government's stated aim is pro-innovation. However, for early stage companies (which are a major source of innovation) who have limited time and money to spend on compliance matters, there may be challenges associated with these companies even just understanding which AI regulators have oversight of their business if further clarity is not provided. A mixture of cross-sector regulators will be difficult for early stage companies to navigate unless further clarity is given.

24. *L2.i. Do you agree that the implementation of our principles through existing legal frameworks will fairly and effectively allocate legal responsibility for AI across the life cycle?*

24.1 Our understanding is that the principles are not to be initially implemented through existing legal frameworks at all, because they are not law.

24.2 There is a risk that the approach taken by regulators – if not streamlined and coordinated - may result in significant barriers to entry, which may unfairly limit AI start-ups (as opposed to larger organisations) from entering the space. This would be to the detriment of those start-ups (a major source of innovation), and in turn, consumers.

24.3 Generally speaking, ensuring that principles are implemented through existing legal frameworks may assist in ensuring that legal responsibility is allocated fairly across an AI's life cycle because liability will be based on an existing body of law. However, in addition to the principles not currently having any legal weight, it is also important to note that liability in the context of an AI's lifecycle is especially complex.

24.4 We have seen recent cases involving potential duties owed by social media platforms to end users for actions performed on their platforms. Likewise, it is currently unclear whether the developer of a foundation model may be responsible for the way in which the model is then applied in practice by end users in circumstances where the developer was aware of the risks the model could present.

24.5 The White Paper envisages engaging a range of experts on the topic of liability. Given that AI tools (both proprietary and open-source) are available to the wider public, these discussions need to take place quickly so as to also equip lawyers and judges with relevant frameworks.

25. *L.2.ii. How could it be improved, if at all?*

Deliberately blank.

26. *L3. If you are a business that develops, uses, or sells AI, how do you currently manage AI risk including through the wider supply chain? How could government support effective AI-related risk management?*

Deliberately blank.

Foundation models and the regulatory framework

27. *F1. What specific challenges will foundation models such as large language models (LLMs) or open-source models pose for regulators trying to determine legal responsibility for AI outcomes?*

27.1 Regarding the principle of transparency and explainability, foundation models (e.g. LLMs) cannot: (i) explain exactly how the technology works; or (ii) predict when a particular input in the form of a prompt will lead to a particular output. There is little that is repeatable in their performance, and evaluations of their output are highly subjective. This makes it difficult for LLMs to comply with the principle of transparency and explainability.

27.2 The operators of LLMs are already gaining significant market dominance. As such, there is unequal bargaining power in the contractual provisions which govern the relationships between those LLM operators and their contractual counterparties (e.g. AI developers). This has the effect of stifling innovation, because those AI developers are squeezed both by the principles, and by the LLM operator.

27.3 Likewise, it is currently unclear whether the developer of a foundation model may be responsible for the way in which the model is then applied in practice by end users, particularly in circumstances where the developer was aware of the risks the model could present.

28. *F2. Do you agree that measuring compute provides a potential tool that could be considered as part of the governance of foundation models?*

Deliberately blank.

29. *F3. Are there other approaches to governing foundation models that would be more effective?*

29.1 There could be requirements for foundation models: (i) to be built on data that has been audited; (ii) which only use data where consent has been given; and/or (iii) where causal models have been in-built. This would essentially mean that even if explainability is less easily attainable, the input data and output can be measured against other principles e.g. fairness, accountability, governance etc.

29.2 See response to *Question F1* above.

Withers LLP, 21 June 2023.

Withers' authors and editors (in alphabetical order): Isaac Black, Henry Farris, Harvey Knight, Richard Stebbing, Giulia Trojano, Ben Williams, Ali Woodcock-West.