

Written Evidence Submitted by Sage (GAI0108)

About Sage

Sage exists to knock down barriers so everyone can thrive, starting with the millions of small and mid-sized businesses served by us, our partners and accountants. Customers trust our finance, HR and payroll software to make work and money flow. By digitising business processes and relationships with customers, suppliers, employees, banks and governments, our digital network connects SMBs, removing friction and delivering insights.

Summary

The application of machine learning and AI to the data that flows across our network will enable us to make predictions about business operations, that help customers gain business advantage.

??

Sage welcomes the Government's important and timely National AI Strategy and a pro-innovation approach to [regulating AI](#). As a global tech company with its headquarters in the North-East of England and as a passionate champion of SMBs, we welcome the Government's ambition for a pro-innovation, regulatory framework for AI in the UK.

This has the potential to unlock the power of AI to help address our shared priorities – increased productivity and sustainable growth – whilst maintaining trust. We support the Government's aspiration to enhance the UK's global reputation as a hub for responsible AI-driven business, whilst continuing to uphold its high standards of regulation.

The links between digitisation and productivity of SMBs are well documented and yet many processes today still require significant manual intervention. AI & ML will be fundamental to delivering Sage's vision to remove barriers, admin burden and friction that SMBs face in growing their businesses and successfully managing their workforce. It would help to knock down barriers for our customers and enable them to thrive.

We believe that now more than ever, with the right AI regulation in place, businesses of all sizes and the government will be able to capitalize on the global opportunities offered by the digital economy.

Sage understands the rationale behind taking a principled and context specific regulatory approach rather than setting an inflexible and prescriptive definition of AI, given its many different applications. However, in the absence of a fully developed standards framework and

compliance guidance, the AI regulatory space is at risk of becoming fragmented and complex when trying to plan for real-world business applications. We need to avoid a situation where businesses or citizens are having to navigate wide disparities and a lack of coherence between regulators.

We believe more consistency is therefore required from Government for different stakeholders. **Sage would prioritise a governance framework that is clear and accessible for all companies, including SMBs.** Sage research reveals that 59% of SMBs¹ found existing data law and regulations too complex to follow an important consideration when developing future AI regulations and solutions. The scale and complexity brought on by different regulatory approaches and possible definitions of AI, could deter deployment by all businesses but especially SMBs who may simply not be sufficiently resourced to assess and take the risk of using AI while remaining compliant. It is important that the regulatory requirements are adequately clear and stable so that SMBs are not deterred from the benefits of AI. It can be difficult for all companies to interpret and incorporate the overlapping regulatory provisions if they are not set out in a regular framework of compliance requirements and guidance.

We therefore urge the UK government to stay engaged in international discussions on AI standards and regulation, which many companies will be brought into compliance with, for example the EU AI Act. By continuing to engage in the international debate, there is not only the opportunity to learn from other markets, but to influence others to take a context specific approach and lead the way.

Today, many large companies and governments are employing AI that can help tackle common challenges.

Therefore, we need the Government to provide a clear standards framework so that all businesses may employ AI on a level playing field.

1. How effective is current governance of AI in the UK?

Effectiveness of the UK approach to AI Governance

1.1 Sage welcomes the ambitions of the UK's pro innovation approach for regulating AI. However, the effectiveness of the approach relies heavily on its interpretation and implementation by the different regulators. A fragmented approach across sectors could lead to an overly complex environment for businesses to navigate. Therefore collaboration and consistency across the regulatory landscape in the UK will be key. We welcome the establishment of the Digital Regulation Cooperation Forum and hope it will take a proactive approach to aligning and building AI principles across sectors.

¹ Sage Edelman poll 1000 US & UK businesses 2019

1.2 The timescales are a complicated issue, as the product lifecycle of various AI products may be several years, and it is important that developers are aware of the regulatory environment in advance. It would be important for regulators to keep up-to-date with the latest applications which may extend beyond purposes taken into account when the regulations were framed.

1.3 However, regulators are much more responsive to shifts in public attitudes and newly elected governments, so there is an amount of expectation management, public education and regulator engagement that may need to be taken account of in the shorter term.

1.4 While the primary focus of this response is on the UK's approach, companies would need to remain cognisant of divergent approaches that are being developed in our other key jurisdictions of their operation, such as EU and US. The divergence and convergence of different regulator regimes will therefore have a major impact on how companies would arrange their operational compliance, along with their strategic judgements about the relative tradeoffs of different business models.

2. What are the current strengths and weaknesses of current arrangements, including for research?

Strengths & weakness of current and proposed AI governance arrangements

2.1 The proposed approach of delegating to individual regulators could potentially lead to uncertainty as judgments are yet to be announced and precedents may be set on a case-by-case basis, which could entail changes in regulation as precedents are set or old rulings are overturned. An unstable and unpredictable regulatory environment will inhibit investment and innovation, so it is important that the UK Government sets out how this could be prevented.

Accessibility for SMBs - standards framework

2.2 Smaller businesses often lack capacity to comply and engage with complex international rules. Introducing a data intermediary which communicates consistent standards, could enable small businesses to benefit from access to wider AI regulatory markets with less risk of breaching regulation.

2.3 As passionate supporters of the UK's small and medium sized business community we are concerned the scale and complexity brought on by different regulatory approaches and possible definitions of AI, could deter the uptake and deployment of AI by businesses without the resources to assess and take the risk of using AI while remaining compliant. It is important that the regulatory burdens are adequately clear and stable so that SMBs are not deterred from the benefits of AI. It can be difficult for all companies to interpret and incorporate the overlapping regulatory provisions if they are not set out in a regular framework of compliance requirements and guidance.

The cost and time needed to comply with unclear or diverging regulatory provisions could act as a barrier – especially to SMBs - slowing down the speed of innovation in the businesses that have the potential to innovate most and contribute to economic growth. A proportionate approach to regulation for SMBs is therefore needed, to avoid an uneven playing field that limits SMB's ability (relative to larger businesses) to leverage AI and innovate at pace. A standards-based framework approach will help clarify the regulatory environment in which AI users of all sizes may operate fairly.

Data sharing arrangements for research purposes

2.4 It would be beneficial to have a framework in place that would enable organizations to use data more freely, subject to appropriate safeguards, for the purpose of training and testing AI responsibly. Further legal clarity is needed on how personal data can be lawfully processed for the purpose of ensuring bias monitoring, detection, and correction in relation to AI systems. Such a clear and accessible framework would provide researchers with the confidence to leverage the many opportunities that AI tooling can enable. This would be extremely beneficial for companies that want to experiment with AI without breaching the regulatory standards and principles. Appropriate safeguards should still be required to help protect the rights of the relevant individuals.

3. What measures could make the use of AI more transparent and explainable to the public?

Ensuring that AI is transparent and explainable, both for compliance and accountability purposes, but also for impact and ethics research.

3.1 It is often the case that the function or direction of an AI's deductions cannot be described with certainty from the outset of a project, and it would be very onerous and disruptive to AI learning models to have to try and explain in real time as the learning process itself was taking place.

3.2 Even if it is possible to identify the purpose of research at the time of AI learning, the nature, structure, scope, and depth of the decisions reached, which may be discovered after collection, often impacts the feasibility of achieving the initially stated purpose.

3.3 Many AI models make deductions from data that are often oblique to human understanding, and it is easier to give post-hoc explanations of its conclusions rather than trying to give real-time or proactive commentary on how it is, or will be, arriving at conclusions.

3.4 GDPR is useful as a model, as it is relatively easy to explain where data came from and what it is being used for. However, if regulators went a level deeper and demanded answers on the AI micro-steps of exactly how decisions were made, it can be easier to explain *how* an AI will learn, rather than *what* it will learn.

3.5. Explainability is not unlike the purpose limitation of GDPR where one can't always foresee the purposes of processing data at the outset, which therefore becomes a major stumbling block before even starting.

3.6 'Black box' models (e.g. a neural network) may often have a better ability to meet a complicated requirement, but come with the trade-off that it is not as straightforward to interpret the reason for a result. More traditional 'white box' models are the opposite - less powerful, more explainable. Clarifying the scope of 'explainability' would give people the confidence to deploy more powerful models in the right situation, without threatening IP. As well as the scope, it's the level of detail required of the explanation that can inhibit innovation. If you could "explain" the model by describing the characteristics of the training data, the steps you have taken to remove bias or other flaws, the testing

you have done on the predictions, and the model type used, this is more realistic than specific predictions. Getting into the detail could be problematic for black box models, as mentioned.

4. How should decisions involving AI be reviewed and scrutinised in both public and private sectors?

Further legal clarity is needed on how personal data can be lawfully processed for the purpose of ensuring bias monitoring, detection and correction in relation to AI systems.

4.1 Sage feels that the Government's proposal that AI regulation should be guided by 'fairness', needs significant work to explain exactly what is meant. 'Fairness' appears to be a loose-fitting concept, and it is likely that over time the definition may shift which is not conducive to the stable regulatory environment that is necessary for efficient development, especially over the timescales with which firms need to plan their AI applications.

4.2 Sage would propose to align the 'fairness' principle with 'human-centred values', which is [the approach taken by the OECD](#) in framing their Values-based principles for AI regulation.' *AI actors should respect the rule of law, human rights and democratic values, throughout the AI system lifecycle. These include freedom, dignity and autonomy, privacy and data protection, non-discrimination and equality, diversity, fairness, social justice, and internationally recognised labour rights. To this end, AI actors should implement mechanisms and safeguards, such as capacity for human determination, that are appropriate to the context and consistent with the state of art.'*

4.3 'Fairness' may mean that companies may need additional oversight measures to evidence that they have made a best effort to avoid unfairness. Sage is sufficiently well-resourced to set up an accountable ethics committee. However, Sage's customer base is primarily SMBs, who often may not have the time or resources to sit and debate its application. If an SMB is developing AI tools, can they be assured they are interpreting the governance framework correctly so it does not inhibit research and development?

4.4 Many SMBs already struggle to interpret GDPR compliance, and this additional regulation (which risks being more complex) should be clarified significantly to avoid it being too burdensome for SMBs to adopt AI products.

4.5 Rather than framing automated processing and decision-making in a regulatory construct akin to "it should not take place, unless these limited grounds allow", it would be preferable from an innovation perspective for the approach to be "it can take place, subject to these transparency and reporting standards".

4.6 The standards could be things like:

- Regular/periodic review, depending on potential impact of technology, human intervention is actively sought and built into automated technologies. For example: reviewing for bias, verifying accuracy, quality assurance.
- Enhanced transparency, For example: explaining how technologies work in plain English, how human involvement is included, how individuals can challenge results etc;

- Enhanced accountability and due diligence. For example: demonstrating how technology developers have built-in fairness by design, demonstrating how fairness/bias is assessed, demonstrating whether non-automated options have been considered etc.

4.7 Without the clarity of a regular standards framework and guidance rubric, organisations may decide not to pursue projects that pose too great a risk to individuals' rights given these wider obligations.

4.8 When considering risk and ethical principles human centred values will place a high priority on personal data. Ethical handling of personal data is critical, and ensuring transparency around how it is handled.

4.9 In the absence of a clear and regular standards framework for the 'fairness' principle, there is a case to be made for holding larger firms to a higher and more rigorous standard than small firms may be able to accommodate if they are not provided the same 'tramlines' as everyone else.

4.10 As passionate supporters of the UK's small and medium sized business community, Sage is concerned that without clarity, there is a real risk that AI "fairness" principles would disadvantage small businesses.

4.11 Transparency will be the key approach for ensuring fairness. There are currently no industry standards for documenting ML datasets, and ML behaviour. This is something the UK can lead the way on, and would help to make AI developers' jobs easier by having a clear framework within which to work.

4.12 The solution to the difficulties of datasets changing, improving and algorithms changing, is to encourage an ethical development process, with a review stage before, during and after the release of AI products. This should involve a documentation and transparency standards, and peer review among experts who can collectively hold each other to account.

4.13 Sage welcomes the UK's ambition to evolve standards as well as regulate, such as expressed in the CDEI's [updated](#) Algorithmic Transparency Standard [published on Github](#). This an example of real leadership on these issues and is to be commended. However, while welcoming the initiative, Sage would advocate for a simpler process to document and evaluate most processes and algorithms. Transparency standards can be higher for higher-risk processes (see **Point 1.5**) but for many low-risk items, the tradeoff between accountability and efficiency can be balanced differently.

5. Are current options for challenging the use of AI adequate and, if not, how can they be improved?

Define legal persons' responsibility for AI Governance:

5.1 Sage would like greater clarity on how far responsibility is to be attributed and in what proportion. For instance, an end user employing a suppliers' proprietary AI product in combination with other AI products and uses them in a way that runs afoul of regulators, then how far up the supply chain does responsibility go? For example the Sage Partner Network offers Independent Software Vendors (ISV) and Developers access to Sage products, developer and educational resources, and technical support. Sage would want to understand if companies developing AI may need to ensure relevant suppliers and

partners would apply the same principles – perhaps by articulating their relationship with Integrated Service Providers, with terms and conditions stipulating how customers use AI products in a compliant manner.

Clarify routes to redress

5.2 It would be important to address how can individuals' rights to redress be maintained when they seek to contest a decision made by AI. This ties into both existing GDPR compliance, but also into the 'explainability' standard the regulators would set. (see points 3.5 - 3.6 above).

5.3 Sage notes the particular restriction in the GDPR relating to AI. The restriction is that individuals have the right not to be subject to decisions based solely on AI (i.e. no human involvement) where there may be a significant impact on them. The restriction does not apply where:

- the decision is necessary for entering into/performing a contract with the individual (likely to be interpreted restrictively); or
- the individual has provided consent to the use of the technology (consent would need to be specific, opt-in, informed and clear/unbundled).

5.4 Human involvement is certainly advisable where there may be significant impact on individuals or entity, as it may be difficult to consistently rely on the conditions above. It would also be worth considering whether significant impact is likely in all cases, and what the test for this would be.

5.5 Sage would suggest that once AI is making decisions which significantly impact entities, transparency frameworks and human oversight are more likely to be needed to document the inherent biases, skews and behaviours in that data. This would provide the necessary quality control and accountability for AI products for any retrospective redresses that may be required.

Regulatory interpretation of principles to challenge AI uses

5.6 This could be complex in areas like "safety" or "fairness" - both for providers of AI and users, especially SMBs without clear guidance. There is a complex interaction with other legislation and regulations: example is using AI for insurance underwriting - already a long-standing debate about the use of data which throws up objective, evidence-based risks, but which can result in certain groups being excluded on grounds of, e.g., age, ethnicity or sexual orientation.

6. How should the use of AI be regulated, and which body or bodies should provide regulatory oversight?

Drawbacks of delegating to regulators

6.1 A disadvantage of the proposed UK approach of context-driven regulation is that rulings may be very contingent, and new regulatory precedents on particular points may unintentionally have sweeping impacts across a variety of other contexts which would be restricted until a new test case is brought to establish the necessary nuances.

6.2 Delegating to different regulators means that AI firms may be burdened by having to engage across a compliance ecosystem involving multiple overlapping areas of regulators' shifting judgements. Broadly it introduces the risk of wide disparities and a lack of coherence across products which may span several regulatory domains.

6.3 With a large SMB client base, Sage is particularly concerned about the impact that this potentially fractured regulatory environment may have on SMBs, which may not have the resources to engage across multiple regulatory environments.

6.4 The input of different regulators is necessary in assessing fairness as AI applications span different domains, and it would be necessary for the regulators to be coordinated. The Digital Regulation Cooperation Forum can play a key role here. The delegation to different regulators may assist in contextual application, but there is also the potential for different interpretations (regulatory guidance and case law) developing in connection with different legislation. This may make it more complex for SMBs to comply where they are having to maintain an awareness of multiple legislation and related interpretations. One of the biggest barriers for SMBs is regulatory fragmentation and having to navigate different regimes.

A standards-based approach to coordinate regulators

6.5 Further consultation and guidance prepared collaboratively with those working in the sector (e.g. AI firms/data scientists/ML engineers) should be encouraged. Ensuring that transparency is one of the key standards could assist in oversight assessments both before, during and after the development of ML products and their applications. A standards-based approach could assist with clarity and help ensure that all businesses are operating within the same framework and on a level playing field.

7. To what extent is the legal framework for the use of AI, especially in making decisions, fit for purpose?

High risk demarcation

7.1 Sage advocates for a regulatory framework that offers protection both to individuals and to companies – especially SMBs. UK data protection regulation is focused on personal data and therefore aims to protect the privacy and rights of individuals and does not protect companies. We would argue that SMBs may be negatively impacted by biased or flawed AI in similar ways to individuals, and humans involved in that business, such as its owners or employees, are also negatively impacted.

Instead of being based solely on the nature of the subject – individual or company – the risk demarcation for AI could be based on the scale of the potential impact. In Open Banking, there are similar risks around the automatic flagging of transactional risk indicators, which can block a payment until greater verification has taken place. For example, AI used to make decisions on companies' creditworthiness could carry the same risks as AI being used for individuals' creditworthiness. On the other hand, AI used by a company to predict their own cash flow could carry a lower risk since the impact of flaws or biases is limited in its scope to the company operating the AI, not third parties, even though the potential negative impact on the company could be high. Also, AI that is used to make inherently low impact predictions, such as which ads to show in social media applications, could be considered low risk regardless of whether they are aimed at individuals or companies. This is an example

where the breadth of the impact is large, but the severity is very low. This approach would be more business-friendly and would follow the model adopted by the PoPIA privacy regulators in South Africa which give protection to both natural persons and “juristic” persons (i.e. companies and organisations).

7.2 Impacts should be demarcated as high risk and should be held to a higher standard of transparency and ethical oversight. This is a general point for AI applications, and there is not yet a legislative basis for consistent distinction, as each regulator would apply existing legislative regulations towards a variety of AI applications in many different regulatory domains. Further consultation is required as to whether an update to legal frameworks would be required for this, or whether an industry standard framework would be sufficient.

7.3 These additional oversight standards would not be necessary for more generic AI purposes where confidential data is not being used, and individuals are not impacted by judgements. This twin-track approach to regulation could allow for innovation not to be impeded for low-impact general AI applications, while ensuring that more high-impact AI applications would be appropriately safeguarded and have to meet higher compliance standards.

8. Is more legislation or better guidance required?

8.1 While Sage welcomes the conceptual foundation of the regulatory system being based on a set of principles, the practical application of the system by regulators could take a standards-based approach. Consolidating the rules and guidance via a coordinator into a standards framework and guidance rubric could reduce the time required (and hence cost) for both SMBs and larger businesses to comply and enable them to use AI for their business more effectively across the different regulators’ domains. Such clarification would be helpful for businesses of all sizes who would otherwise struggle to interpret the potentially fractured regulatory landscape of the AI space.

8.2 Sage discusses the a standards framework and guidance rubric elsewhere in the submission, in particular the case for their introduction, the important accessibility benefits it would bring to SMBs, how the OECD ‘human-centred values’ are a helpful example that avoid the risk of ambiguity of the UK position and some practical suggestions of particular standards guidance that would be helpful.

9. What lessons, if any, can the UK learn from other countries on AI governance?

As far as we are aware there do not appear to be off-the-shelf models that can help immediately with improving the effectiveness of UK AI regulation.

9.1 Sage wants the UK to be at the forefront of ethical, innovation-led AI regulation and supports the Government’s ambitions to allow the governance structures to be innovation and research-led.

9.2 The reason for this is that ‘AI’ is a catch-all term for a diverse and rapidly developing field. There will be many datasets available that might be useful for assessing specific problems around regulatory

effectiveness. There are early individual examples which can be referenced, but they are in academic ML/AI ethics research journals. For instance, the Distributed AI Research (DAIR) Institute, which documents the impact of AI on marginalized groups. However, in order to access the datasets, regulators would have to dig through the papers and have a high-level appreciation of the cutting edge of academic literature in this area.

9.3 The problem with grafting any lessons from those datasets onto others, is that the use-cases are all very specific.

9.4 For example the DAIR Institute's founder Timnit Gebru, whose research covers on algorithmic bias in data mining, focusing on the biases learned by large language models (LLMs), an AI approach that was only invented recently. Gebru's work quantifies outcomes and biases in LLM datasets, and proving for instance, that the AI has learned to associate women's names with traditionally female occupations, or African American names not being associated with white-collar professions. In other words datasets do exist, and can be referenced, but only for those problems, and solely with LLMs – such datasets may often not be relevant to other AI products which were set up to solve different problems on different datasets.

9.5 There will be many other ML problems throughout the world that don't have those datasets, or are just as specific in their own application, but where less ethical research has been done.

Relevant datasets will require consistent documentation, transparency and review standards.

9.6 The solution to the difficulties of datasets changing, improving and algorithms changing, is to encourage an ethical development process, with a review stage before, during and after the release of AI products. This should involve a documentation and transparency standards, and peer review among experts who can collectively hold each other to account.

The Government could fund projects to ensure data sources are domestically produced and relevant to the particular DEI challenges of the UK, rather than importing e.g. US fixes for US DEI issues.

9.7 While Sage works across many markets, as a UK-based business, Sage would also note the importance of ensuring domestically-produced UK datasets for the purpose of transparent ethical monitoring for both individual regulators and at system level. It is important that this supply is regularly refreshed and recalibrated, in order to ensure that it is up-to-date with the latest developments in AI research and applications, which are a fast-moving field. This is perhaps an area of research which the UK Government – through its agencies such as UK Research & Innovation and its subsidiary grantmaking councils – should consider sponsoring as a long-term workstream, in order to assist the UK AI industry to have the most relevant datasets appropriate for a UK societal context.

9.8 There is always a risk that - as the largest volume of AI ethics research is conducted in the United States US-derived research and datasets may be shared as best practice, and the particular fixes of the US' structural social biases are exported to other countries, regardless of the relevance to the other countries' social texture. Diversity, Equity and Inclusion discussions which happen in – for instance –

the United States, the UK or France, all take place in very different social-historical textures, and it is important that the datasets are relevant and sensitive to the particular domestic challengers of each. US ethical AI datasets aim to correct for US DEI issues, and are not appropriate for a direct graft into a UK context, although they have great value as an example of how to test in general for particular biases and fairness.

9.9 The UK has a responsibility to ensure domestic dataset projects are attuned to UK needs and sensitive to the particular DEI texture here. Further discussions on how the UK Government and its agencies could support this research would be very welcome, and would contribute to economic growth and UK regulatory leadership, as the UK AI industry can be confident that it is operating with ethical frameworks which are well-calibrated to a UK context both societally and in terms of regulator's concerns.

Intent to export standards across the world and leading international cooperation

9.10 Sage would encourage the Government to make full use of international forums such as the G7 and G20 to drive the AI regulatory agenda and create common standards from the UK's proposed model.

9.11 Sage welcomes the government's ambition to be a 'rule-setter' but would like to understand more about where the UK intends to lead with dominant powers like China, the EU and US looking to go their own way.

9.12 Sage have substantial EU & US operations, and need to look across our international markets. The outcomes of AI governance reforms in any of these markets could have a major impact on Sage's commercial products,.

9.13 Sage is UK-headquartered and thereby has a particular responsibility to provide industry leadership in this area. From a domestic growth perspective of supporting UK leadership in the global AI space, it is important that the Government gets this right and Sage is very supportive of this ambition, provided that accessibility for SMBs is prioritised.

9.14 While the UK's Alan Turing Institute is setting up an initiative to lead in shaping global technical standards for AI, alongside the Centre for Data Ethics and Innovation, it would be a useful point of comparison to monitor the work around AI risk management and standards of the American National Institute of Standards and Technology. Notably, they are taking a guidelines, rather than regulation-based approach, and are at an early stage of trying to work toward some international standards.

9.16 The UK's intention to be a regulatory leader and exporter is very positive and useful goal especially as many markets are behind the UK's position. To make this more than an ambition, however, they need a plan from Government for how this will be achieved. Sage sees especially positive opportunities for UK tech to sell / operate AI in US and EU markets. Sage would propose that this is made a priority.

(November 2022)