

Principles and power: The ODI members' roundtable on the AI White Paper

To inform our response to the government's White Paper, [AI regulation: a pro-innovation approach](#), the ODI hosted a virtual roundtable for its members on Tuesday 13 June 2023.

Attendees from a range of organisations including regulators, big tech, small tech, SMEs and civil society groups participated in the event, where we were able to gather a range of insights and reactions to the government's approach, as well as ideas of missed opportunities, and concerns about risks that had been overlooked. We based the different sessions of the roundtable on several - though not all - of the questions from the consultation, as well as giving space for more general reflections.

This short paper summarises the discussions at the roundtable. The ODI's full response to the White Paper is [here](#) and we've written a blog summarising our thoughts on the White Paper [here](#).

Key points included:

- There was an even spread of views on whether the government's general approach (principles-based, contextual, sector-specific, no overarching regulator) was the right one. Scepticism focused on regulatory complexity, confusion and a lack of regulation, a lack of new powers and a lack of focus on the data powering AI models; more positive comments thought focusing on the use of AI in context in particular sectors was sensible, but also worried about coordination
- There was broad support for the government's proposed principles (safety, security and robustness; transparency and explainability; fairness; accountability and governance; contestability and redress), though people wondered where considerations about data should fit. They also felt real-world harms to people and society were being missed, including environmental harms and risks to marginalised groups, as was the speed of change and the international dimension to debates about AI.
- Members thought the government approach broadly struck the right balance between innovation and risk, though there were worries that future-proofing seems impossible and not enough has been learned from the last few decades of digital growth and that there isn't enough certainty (around direction, guardrails and legitimacy) provided for businesses to innovate.

Introduction

In our introduction to the roundtable, we noted that the government wants the UK to be ‘the best place in the world to build, test and use AI technology’.

According to the government’s [Science and Technology framework](#), AI is one of five critical technologies - alongside Engineering Biology; Future Telecommunications; Semiconductors; and Quantum Technologies -, and the UK is already well-placed on AI: the White Paper notes that the UK is ranked third in the world for AI research and development,¹ and is home to one-third of all of Europe’s AI companies.

We noted that the government’s approach was, broadly, regulation-based on cross-sector principles without creating a new regulator - existing regulators would apply the principles in their own domains. There would be no new regulator, but there would be ‘central functions’ to support regulators and monitor and evaluate the effectiveness of the framework. The regulatory system would be ‘deliberately agile and iterative’.

For a summary of the government’s proposed approach, jump to the [appendix](#).

This document will discuss:

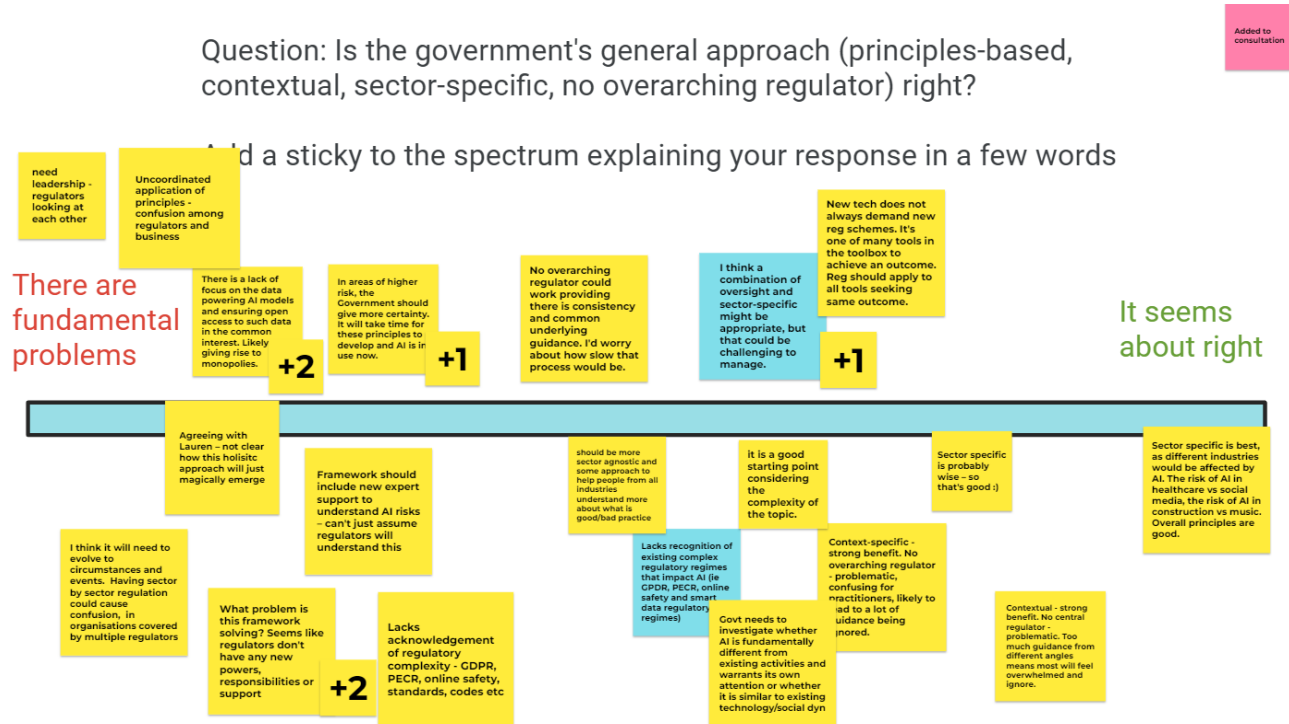
Views on the government’s general approach	3
Views on the principles	5
Views on striking the right balance	8
Views on whether the right sectors and needs were represented	10
Breakout groups	12
Group 1: Transparency	12
Group 2: Central function and statutory requirements	13
Group 3: Regulation and avoiding duplication of roles	15
Group 4: Does the approach work for all/cover the right risks	16
Views on innovation and increasing trust in AI	18
Final comments	20

¹ The latest edition of [the Tortoise AI Index](#), published after the consultation closed, sees the UK slip to fourth thanks to Singapore’s rapid rise up the rankings.

Views on the government's general approach

We asked: **Is the government's general approach - principles-based, contextual, sector-specific, no overarching regulator - the right one?**

Participants were invited to add their thoughts on a Jamboard, along a spectrum - reflecting 'Fundamental problems' on the left, and 'seems about right' on the right.



There was a relatively even spread of views. Concerns, those on the left - towards 'fundamental problems', starting with the furthest left - tended to focus on the difficulties of coordination between regulators, particularly without additional powers and support:

- *Need leadership – [rather than] regulators looking at each other.*
- *I think it will need to evolve to circumstances and events. Having sector-by-sector regulation could cause confusion, in organisations covered by multiple regulators.*
- *Uncoordinated application of principles - confusion among regulators and business.*
- *Agreeing with [a speaker during the roundtable] – not clear how this holistic approach will just magically emerge.*
- *There is a lack of focus on the data powering AI models and ensuring open access to such data in the common interest. Likely giving rise to monopolies. [+2]*
- *What problem is this framework solving? Seems like regulators don't have any new powers, responsibilities or support. [+2]*

- *Framework should include new expert support to understand AI risks – can't just assume regulators will understand this.*
- *In areas of higher risk, the Government should give more certainty. It will take time for these principles to develop and AI is in use now. [+1]*
- *Lacks acknowledgement of regulatory complexity – GDPR, PECR, online safety, standards, codes etc.*

Those feeling more positively about the framework -starting with the furthest right - tended to acknowledge the challenges that would come with the approach, however desirable:

- *Sector-specific is best, as different industries would be affected by AI. The risk of AI in healthcare vs social media, the risk of AI in construction vs music. Overall principles are good.*
- *Sector-specific is probably wise – so that's good :)*
- *New tech does not always demand new reg schemes. It's one of many tools in the toolbox to achieve an outcome. Reg should apply to all tools seeking same outcome.*
- *Context-specific - strong benefit. No overarching regulator - problematic, confusing for practitioners, likely to lead to a lot of guidance being ignored.*
- *It is a good starting point considering the complexity of the topic.*
- *I think a combination of oversight and sector-specific might be appropriate, but that could be challenging to manage. [+1]*

Comments in the middle of the spectrum acknowledged some of the possible advantages, but also the challenges:

- *The Government needs to investigate whether AI is fundamentally different from existing activities and warrants its own attention or whether it is similar to existing technology/social dynamics.*
- *Lacks recognition of existing complex regulatory regimes that impact AI (ie GDPR, PECR, online safety and smart data regulatory regimes).*
- *should be more sector agnostic and some approach to help people from all industries understand more about what is good/bad practice.*
- *No overarching regulator could work providing there is consistency and common underlying guidance. I'd worry about how slow that process would be.*

Views on the principles

Our second question was **whether the five principles identified by the government in the White Paper were the right ones.**

One participant said that they “generally think these are good principles but oversight is key.”

The five principles, and some specific comments on them individually, are:

- Safety, security and robustness.
- Appropriate transparency and explainability:
 - *Who in the AI lifecycle is responsible for transparency and explainability?*
 - *Transparency is not just about what data but how different datasets are being joined together.*
- Fairness
 - *I think [the principles are] about right although fairness is quite loose. [+2]*
 - *Is equity a better term than fairness?*
 - *Is access part of fairness? Feels like it would be part of equity.*
- Accountability and governance
- Contestability and redress:
 - *You don't usually hear about contestability and redress, so I'm pleased that this is in here[Contestability and redress] is really key. I think this demands a stronger centralised response.*

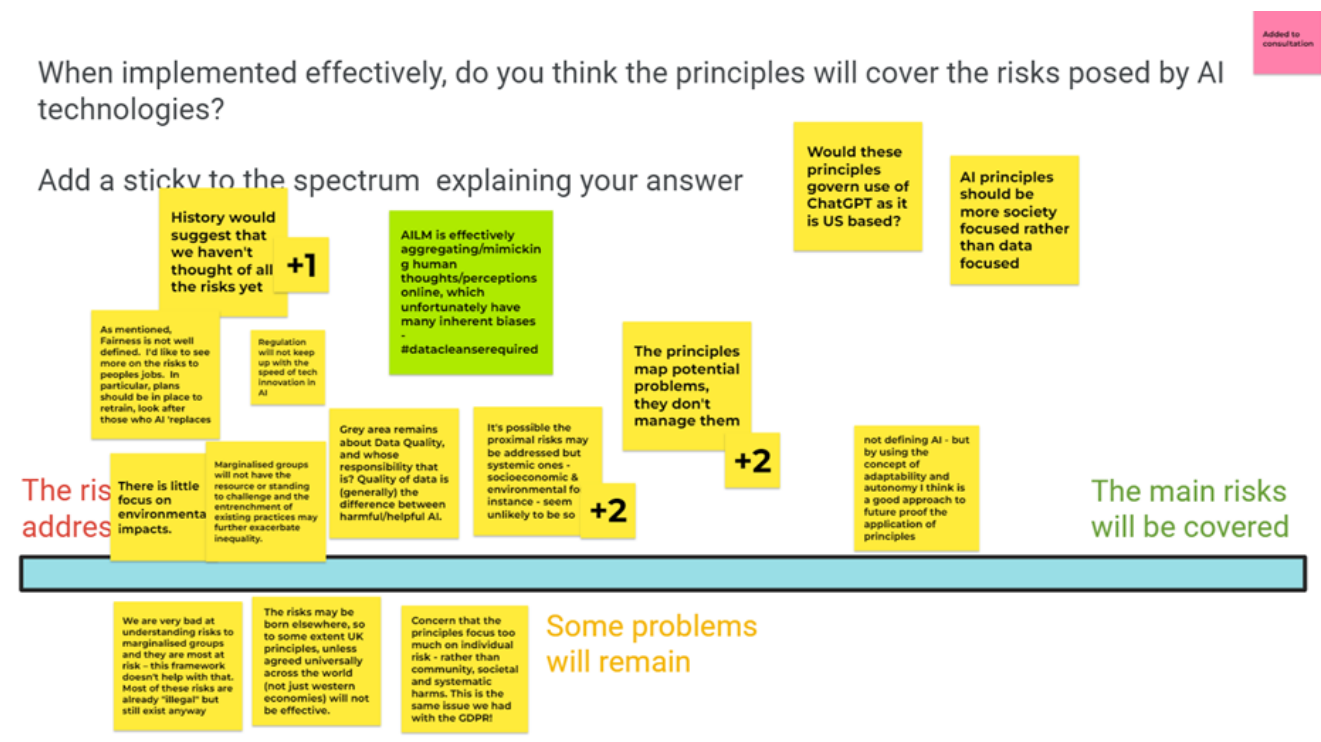
There were several comments about where considerations about data would fit:

- *Obviously key to AI is the data set. This needs to be rightly reflective of what is being modelled. I'm not sure where that lands here - fairness? I feel it is broader than just fairness.*
- *Access to data! If we want to support AI innovation we need accessible high quality public datasets. Regulators should have a role in this! Hopefully a popular opinion!*
- *Agreed on Data access, but it's more than just that - [5 Safes framework](#) could be helpful - nature of data, use of data, people impacted by data, setting to use data and outputs. [A second comment wondered if the 5 Safes 'could also be complementary or helpful'].*
- *Where does mitigating AI from aggregating existing bias fit into the principles?*
- *Not much to address power in markets gained by accumulating data.*
- *Principle relating to 'Data access and Usage'?*

And there were a few other comments:

- If AI is different from other technology it's because of the social implications. So the principles should be more society-centric than data/technology centric. [+1]
 - Yes, so not just about potential harm to individuals, but also potential harm to wider society.
- I think environmental should be explicitly separate.
- They are quite "low level" principles. Not talking about outcomes for real people. [+1]

We asked a second question about the principles: **When implemented effectively, do you think the principles will cover the risks posed by AI technologies?**



Participants were asked to add their comments along a spectrum – from ‘The risks will not be addressed’ on the left, through ‘Some problems will remain’, to ‘The main risks will be covered’ on the right.

Most of the comments were made towards the left hand side of the spectrum. The only particularly positive comment was:

- *Not defining AI - but by using the concept of adaptability and autonomy I think is a good approach to future proof the application of principles.*

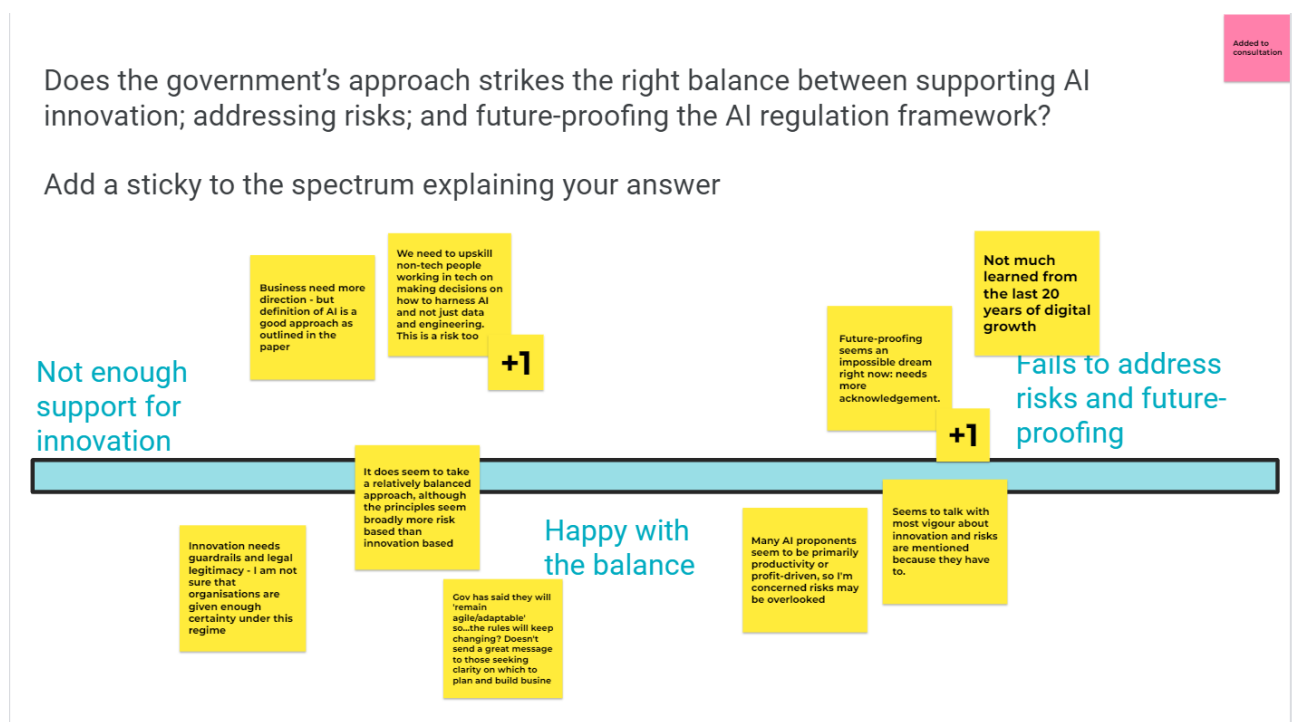
Concerns included:

- *Major societal risks being missed:*
 - *AI principles should be more society focused rather than data focused*
 - *It's possible the proximal risks may be addressed but systemic ones - socioeconomic & environmental for instance - seem unlikely to be so [+2]*
 - *There is little focus on environmental impacts*
 - *Concern that the principles focus too much on individual risk - rather than community, societal and systematic harms. This is the same issue we had with the GDPR!*
- *Risks to marginalised groups:*
 - *We are very bad at understanding risks to marginalised groups and they are most at risk – this framework doesn't help with that. Most of these risks are already "illegal" but still exist anyway*
 - *Marginalised groups will not have the resource or standing to challenge and the entrenchment of existing practices may further exacerbate inequality.*
- *The speed of change:*
 - *History would suggest that we haven't thought of all the risks yet [+1]*
 - *Regulation will not keep up with the speed of tech innovation in AI.*
- *The international dimension:*
 - *Would these principles govern use of ChatGPT as it is US based?*
 - *The risks may be born elsewhere, so to some extent UK principles, unless agreed universally across the world (not just western economies) will not be effective.*
- *Some more general points, some picking up on the concerns about the principles in general:*
 - *The principles map potential problems, they don't manage them [+2]*
 - *AI language models are effectively aggregating/mimicking human thoughts/perceptions online, which unfortunately have many inherent biases #datacleanserequired*
 - *Grey area remains about Data Quality, and whose responsibility that is? Quality of data is (generally) the difference between harmful/helpful AI*
 - *As mentioned, Fairness is not well defined. I'd like to see more on the risks to people's jobs. In particular, plans should be in place to retrain, look after those who AI 'replaces'.*

Views on striking the right balance

Our next question was: **Does the government's approach strike the right balance between supporting AI innovation; addressing risks, and future-proofing the AI regulation framework?**

Again, we asked participants to place their thoughts on a spectrum, ranging from 'not enough support on innovation' on the left to 'fails to address risks and future-proofing' on the right.



There were some concerns on both ends of the spectrum. Those worried that there wasn't enough on addressing risk and future-proofing said:

- *Future-proofing seems an impossible dream right now: needs more acknowledgement. [+1]*
- *Not much learned from the last 20 years of digital growth.*
- *Seems to talk with most vigour about innovation and risks are mentioned because they have to.*
- *Many AI proponents seem to be primarily productivity or profit-driven, so I'm concerned risks may be overlooked.*

Those worried that there wasn't enough support for innovation said:

- *Innovation needs guardrails and legal legitimacy - I am not sure that organisations are given enough certainty under this regime.*
- *Business need more direction - but definition of AI is a good approach as outlined in the paper.*
- *It does seem to take a relatively balanced approach, although the principles seem broadly more risk based than innovation based.*
- *We need to upskill non-tech people working in tech on making decisions on how to harness AI and not just data and engineering. This is a risk too. [+1]*
- *Government has said they will 'remain agile/adaptable' so...the rules will keep changing? Doesn't send a great message to those seeking clarity on which to plan and build businesses.*

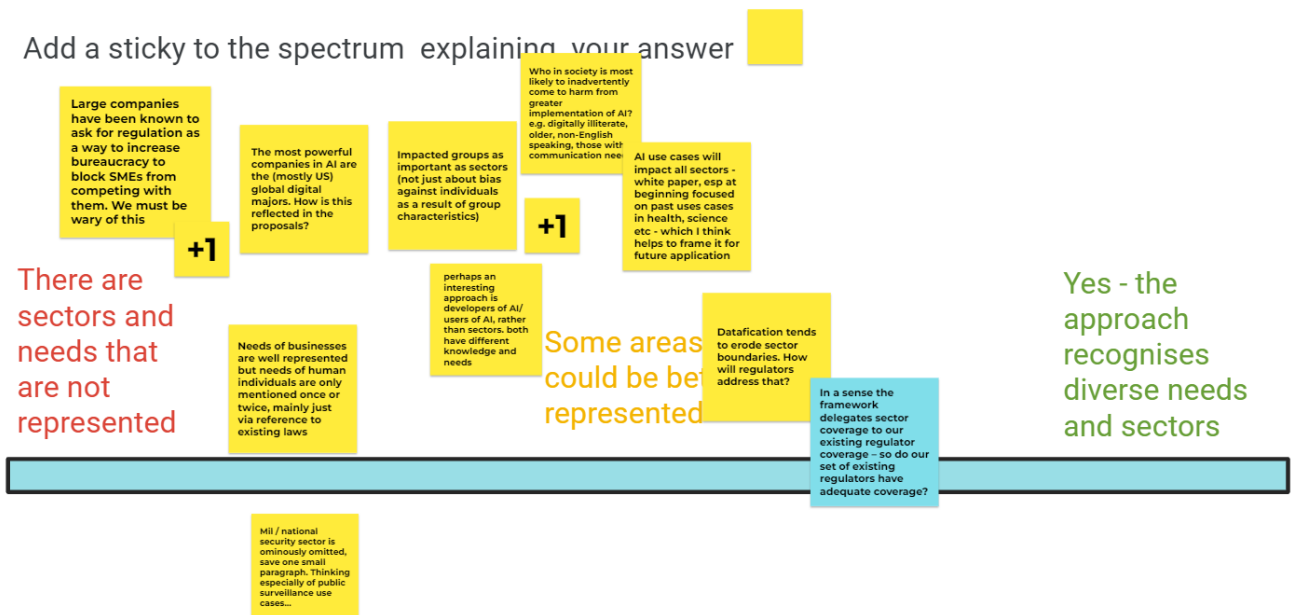
Views on whether the right sectors and needs were represented

We asked: **Are all sectors and needs represented in the approach? If not, what's missing?**

We again asked participants to put their thoughts on a spectrum, from there being sectors and needs not represented on the left, through 'some areas could be better represented', to the approach recognising diverse needs and sectors.

Are all sectors and needs represented in the approach? If not, what's missing?

Add a sticky to the spectrum explaining your answer



Most comments tended towards the left hand side, that there were sectors and needs that could be better represented.

- *There were concerns about those not properly considered:*
 - *Needs of businesses are well represented but needs of human individuals are only mentioned once or twice, mainly just via reference to existing laws*
 - *Impacted groups as important as sectors (not just about bias against individuals as a result of group characteristics) [+1]*
 - *Who in society is most likely to inadvertently come to harm from greater implementation of AI? e.g. digitally illiterate, older, non-English speaking, those with communication needs*

- *Perhaps an interesting approach is developers of AI/ users of AI, rather than sectors. both have different knowledge and needs*
- *Military/national security sector is ominously omitted, save one small paragraph. Thinking especially of public surveillance use cases.*

There were two different points about tech giants:

- *Large companies have been known to ask for regulation as a way to increase bureaucracy to block SMEs from competing with them. We must be wary of this. [+1]*
- *The most powerful companies in AI are the (mostly US) global digital majors. How is this reflected in the proposals?*

One participant appreciated the White Paper's approach:

- *AI use cases will impact all sectors - white paper, especially at beginning focused on past uses cases in health, science etc - which I think helps to frame it for future application.*

Others had questions about the practicality of the regime for regulators:

- *Datafication tends to erode sector boundaries. How will regulators address that?*
- *In a sense the framework delegates sector coverage to our existing regulator coverage – so do our set of existing regulators have adequate coverage?*

Breakout groups

Participants then split into four groups, each discussing and using a Jamboard for several linked questions.

Group 1: Transparency

The first question was: **Would requiring organisations to make it clear when they are using AI adequately ensure transparency?**

There were some high-level comments about whether transparency was the right concept, with one participant noting “The user need is ‘I don't want to be discriminated against’ – transparency is part of that but not necessarily sufficient.” Another wondered if transparency was “Less about AI, perhaps more about decision-making in general. But AI makes decisions more inscrutable.”

There were several comments about the specific transparency requirements and how people would understand them:

- *It seems so, but that would depend on whether that disclosure is buried in an end user agreement (i.e. those things no one reads!).*
- *Also how/what for exactly.*
- *Transparency on data quality is more important than transparency on the presence/or not of an algorithm. We need to get good at explaining specific concepts in layman's terms. Just saying 'FYI this software makes use of AI technologies' would do more harm than good. Another participant said we should allow people the opportunity to understand AI as well as they can, rather than jumping straight to simplification.*
- *How are individuals meant to interpret use of AI?*

The second question was: **What other transparency measures would be appropriate, if any?**

Suggestions included:

- *Proof that AI isn't biased.*
- *Referring back to the earlier point about understanding, ‘demystifying AI would be helpful’.*
- *Logos or branding could be used to alert users to where AI was being used.*
- *Using existing structures – specifically, the Privacy and Consumer Advisory Group – to play an advisory role and think about the ethical implications of AI.*

The third question was: **How far can transparency build public trust?**

Answers here were:

- *Again building on previous points about understanding and explainability, need more than just transparency – education, proof of non-bias, supporting structures to challenge, etc.*
- *Transparency could involve stating the level of AI being used in an information transaction. Some businesses claim to use AI to appear innovative but it's very rudimentary.*
- *More “good news” AI news stories would help.*

The final question was: **Beyond measures already discussed, what would help us to evaluate the impact of AI?**

There were two concrete suggestions:

- *Tracking the narrative of biases online over time? Will AI positively or negatively affect the level of bias?*
- *Where has AI been shown to reduce inequality? Or other examples of where AI has actually mitigated one of the risks we usually associate with it.*

Group 2: Central function and statutory requirements

The first question for this group was: **Is introducing a legal duty on regulators to have due regard to the principles enough to manage risks while being flexible enough for innovation?**

Several of these built on concerns from earlier answers:

- For the moment, maybe, but it feels like this landscape is moving so quickly. One participant asked the rhetorical question: has regulation so far kept pace with technology? Even a subject like vaping shows where regulation has not kept pace.
- No recognition of power deriving from accumulating data. Should regulators have to assess that? Could be, but even within that, a lot depends on the strength of the legal obligation.
- The principles would be more convincing if based on a clear admission of the kind of harms and inequities that come from using data.
- We already have over 100 laws that regulate data in the UK, then we would have to overlay every AI regulation?!
- People want regulators to represent them robustly so regulators need to be compelled to do what is implicitly expected of them.

One participant, who had worked with a regulator, was concerned that a regulator having to have ‘due regard’ would be in addition to their legal duties – an overloaded regulator would pay more attention to the latter, so having regard to guidance would be too weak.

The second question followed the first, asking: **If not, what would be more effective?**

One comment was for the need to ‘look at the underlying principles, rather than principles specific to tools or technologies.’ These could include ‘do no harm’ or avoiding ‘negligence’. Another participant noted the ‘precautionary principle’ at play in the judicial system, and the need to look at possible harms before deploying new systems.

One participant noted the importance of having fundamental principles before taking a position, and asking what the rights, responsibilities and harms were. The Five Safes framework was again cited. The participant observed that regulation had become very reactive, rather than starting with the principles we want to run society by.

Fines might be an important part of the system of regulation, but one participant noted that fines were not enough – they were seen as the cost of doing business for too many large companies. One participant also felt competition doesn’t have a great record” in helping to regulate digital markets.

There were also general worries about the funding and capabilities of regulators, and a sense that much of regulation and reform was designed to reassure businesses rather than the public and was not as open to recognising the harms from datafication as they should be. There might be greater credibility if governments admitted the challenges of regulating a new technology and acknowledged past mistakes, though this was difficult.

The final two questions covered the proposed central functions: **What central functions should there be, and what, practically, do they look like? What, if anything, is missing from the central functions?**

Two points were raised here. The first was about AI sovereignty in the UK – a lot of data storage, especially with the move to the cloud, is with companies not based in the UK. The second was a “need to acknowledge that there will be unknown unknowns. Stop worrying about regulatory certainty: [it’s] unfeasible at the general long-term level.”

Group 3: Regulation and avoiding duplication of roles

This group's first question was: **What can the government do to reduce overlapping, duplicative or contradictory guidance on AI issued by different regulators?**

There were worries about the challenges, which might include “multiple fines from different regulators...a lack of coordination, a lack of coordinated guidance”, which a ‘working group or board of regulators’ might help with. One participant suggested the Financial Conduct Authority’s Joint Regulatory Oversight Committee might be a model for a working group or board. Another participant noted that “business [is] still not sure how to internally structure compliance function – data protection [is] taking the lead here.”

For other participants, this might not be a bad thing, wondering if it would be ‘acceptable’ for “one regulator, especially ICO, [to] take the lead”, and put forward the view that “regulators have very clear remits – [it] should make sense to apply principles and then have the context-specific application. ICO has a broad overview – so [the] balance is there.” The existence of relevant regulators meant that one participant thought new guidance could be limited, as aspects could be incorporated into existing remits.

The second question was: **Would the regulators currently in place cover all the risks posed by AI? If not, why not?**

One participant was positive about the role of the central functions here, thinking they could “play an important role regarding risk registers, etc”, though “coordination and coherence [need] to be thought of – as well as applications of AI.” They thought there would be concerns about risks – another agreed, saying they are still not known, but the white paper’s definition of AI was “a good start.” Another wondered: “Will people want others to take the lead? Who covers the remit?”

The third question was: **Which non-regulatory tools for trustworthy AI would be most helpful in achieving the government’s principles?**

There was an acknowledgement that “auditing AI is a really new field – lots of how we approach is from a question base and the mechanics of that. [There is a] Need for government and civil society to understand these questions.” Organisations might also be ‘fragmented’, as different teams use AI in different ways.

For several participants, data standards “would be a good starting point”:

- *How do we apply [those] and standardise [the] use of data in an organisation?*
- *How can using [the] principles help the creation of industry bodies on a practical level to implement standards on use of AI?*
- *[The Alan] Turing [Institute] has done interesting work linking standards to principles. Useful for tackling lack of coordination and that framework is non statutory. Useful approach from an AI standpoint.*

It would be important to engage stakeholders at a technical level in thinking about what the standards would be – and their interoperability.

The final question was: **What should the policy position on IP law/data replication and generative AI be?**

Participants thought the role of IP law and contractual relationships around the use of data needed to be debated further, in context and ‘holistically’ given it ‘fuels the technology’ around AI.

Group 4: Does the approach work for all/cover the right risks

The first question to this group was: **What challenges are likely to arise from applying principles across different AI applications and systems?**

The replies were:

- *Where can I use my AI models? If trained on UK data, can I use the model in other countries? Compatibility across countries is vital for us. If different sectors have different rules there's an increased risk of misinterpretation.*

The second question was: **What could the government do to recognise the differences in risk, and requirement, of the public and private sectors?**

There were some sector-based responses:

- *Grouping the possible impact on sectors. The risk of AI in healthcare vs social media, the risk of AI in construction vs music. Overall principles are good.*
- *The differences in sectors should mirror the differing privacy requirements of types of data, so that they do not clash. i.e use the same sector/data categories.*

Others answers were more general:

- *Businesses need to have a mitigation strategy for all AI products e.g. handover to a human when the chatbot fails.*
- *Representative Datasets - to avoid biases. Guidance for these.’ [+1]*
- *Tackle the data being collected on us by big tech who them have a monopoly on this data.*

And one participant said: “move more quickly! AI is being trained now.”

The third question was: **How could the government support effective AI-related risk management?**

The answers were:

- *Some required transparency around measuring and managing bias and risks.*
- *Guidance based on categories of sectors and the impact they could have, e.g. need for representative datasets in healthcare to avoid biases.*
- *Acknowledge severe impacts may come in many sectors – so probably need the kind of robustness we usually see in narrow areas like medicine applied more broadly.*

The final question was: **How do businesses manage AI risk including through the wider supply chain?**

Participants felt strongly that “AI risk and governance needs to be an explicit discipline in organisations”, that an AI ethics team should be consulted, and that certification - perhaps kite marks - for systems and suppliers should be rolled out, and quickly.

Views on innovation and increasing trust in AI

We brought the whole group back together to consider several questions around innovation and increasing trust in AI.

When asked if there were **additional activities that would help businesses confidently innovate and use AI technologies**, the answers were:

- More cognitive diversity across the AI lifecycle.
- Standards for governance and accountability.
- Sandbox environments.

One participant questioned if, given that “AI business is already exploding – does it actually need more help? Shouldn't government be focusing on the things that private sector won't – e.g. risks?”

The **existing barriers to innovate and use AI** cited were:

- *Data organisation.*
- *Culture.*
- *Limited understanding of AI [by] elected officials or other key decision makers.*
- *Restrictions in relation to the access and use of data.*

One participant thought data privacy was a barrier: “not saying it shouldn't be, though.”

Two comments focused on the machine and human readability of data:

- *Data designed to be machine readable should be more human friendly for non-techies to engage with it, before deciding how to use it for AI.*
- *This is not a data literacy issue. Bad UX means bad design and data is designed to be consumed by machines not humans.*

Additional activities that would help individuals and consumers confidently use AI technologies included:

- *Innovation labs and roadshows to demystify AI across a range of sectors.*
- *Cultural awareness and ethics for AI. National or global data aggregation of data by AI could marginalise certain cultures.*
- *People would have more confidence in government if government talks more about innovation being for public good, not any innovation being good in itself.*
- *Showing some proof that AI risks can be effectively addressed [e.g. a badge of approval].*

One participant questioned whether “consumers even know they are using AI technologies today or have a choice in the future?”

When asked, **Should these activities be delivered by the government, regulators or a different organisation?** our participants responded:

- *A harmonious combination of public and private regulatory partnership would be ideal... albeit quite optimistic.*
- *Regulatory guidance in the absence of regulatory framework – principles are a start.*
- *Must be independent – needs to hold government to account too.*

Final comments

Asked for any **final thoughts on the government's overall approach**, our roundtable participants said:

- How this approach differs from the EU approach and their AI Act - which might end up being the gold standard - will be interesting, especially as white paper talks about interoperability.
 - Another participant added “and Australia, NZ, USA, etc.”
- The Linked Organisation of Local Authorities (LOLA) would be interested in a unified approach to AI.
- Bias is presumed to be a bad thing – don't think it necessarily is. Sometimes we want to build in bias to fix a problem.

As for **missed opportunities, flaws, or gaps in the government's approach**:

- OpenAI/GPT has already won through being 100x faster than government at innovating – how can we compete against that?
- ChatGPT is so dominant already and is largely unchecked – only heavy regulatory/statutory action will help. Unclear if this framework will have any impact.
- Business-to-business and government-to-business data sharing – equivalent to EU Data Act – UK's approach to this and how it relates to AI [+1]
- AI is a global technology, so has the UK government taken a global perspective to their AI approach, or a UK-specific approach?

Other **final observations, and other comments** from throughout the roundtable, included:

- Would like to hear more about WHY we've picked these principles? Is it because of what's gone wrong over the last 20 years? If so, just say so! [+1]
- Different countries may come to different findings on copyright because different judges will apply their laws to it differently.
- Copyright law in 2011 wasn't designed to be applied to data. Imbalances continuously growing between companies and between companies and consumers. Nothing addresses that.
- Want to see more talk about 'biased data', inc. defining the term 'bias'. Much is discussed about how to reduce or 'remove' bias, but that's not always desirable.
- We need better understanding of uptake across sectors, and AI should not just be about the unicorns: “the innovation department could talk about how it improves people’s lives not just make people richer”, meaning the need for an understanding of the real public benefit.
- There was some discussion of data privacy, some feeling that in still getting to grips with it the initial reaction has been tighter regulation organisationally out of risk aversion, others suggesting that privacy is likely more of a barrier now than it will be in the future, with people on that curve of getting better at managing privacy.

- Data is designed to be machine readable, but not human readable – have we ever addressed the challenge of the human side of data?

We hope you have found reading about the ODI members' roundtable as fascinating as we did to convene it. The ODI will be hosting regular roundtables for our members, covering a broad range of topics, from consultations, to key manifesto asks, to Q&As with industry leaders and policy makers. Find out more about ODI membership [here](#).

Appendix

A bullet point summary of the White Paper

The **five principles** would not, initially, be statutory, but there was an intention over time to introduce a statutory duty on regulators to have ‘due regard’ to them. The principles are:

- Safety, security and robustness.
- Appropriate transparency and explainability.
- Fairness.
- Accountability and governance.
- Contestability and redress.

The **seven proposed central functions** supporting regulators in their efforts are:

- Monitoring, assessment and feedback.
- Support coherent implementation of the principles.
- Cross-sectoral risk assessment.
- Support for innovators (including testbeds and sandboxes).
- Education and awareness.
- Horizon scanning.
- Ensure interoperability with international regulatory frameworks.

The government has stated **three objectives** for the regulatory regime:

- Drive growth and prosperity by making responsible innovation easier and reducing regulatory uncertainty.
- Increase public trust in AI by addressing risks and protecting our fundamental values.
- Strengthen the UK’s position as a global leader in AI.

It has also outlined **six characteristics** for the regime:

- Pro-innovation: enabling rather than stifling responsible innovation.
- Proportionate: avoiding unnecessary or disproportionate burdens for businesses and regulators.
- Trustworthy: addressing real risks and fostering public trust in AI in order to promote and encourage its uptake.
- Adaptable: enabling us to adapt quickly and effectively to keep pace with emergent opportunities and risks as AI technologies evolve.
- Clear: making it easy for actors in the AI life cycle, including businesses using AI, to know what the rules are, who they apply to, who enforces them, and how to comply with them.

- Collaborative: encouraging government, regulators, and industry to work together to facilitate AI innovation, build trust and ensure that the voice of the public is heard and considered.

And **four elements** for the regulatory framework:

- Defining AI based on its unique characteristics to support regulator coordination.
- Adopting a context-specific approach.
- Providing a set of cross-sectoral principles to guide regulator responses to AI risks and opportunities:
 - The principles clarify government's expectations for responsible AI and describe good governance at all stages of the AI life cycle
 - Their application will initially be at the discretion of the regulators, allowing prioritisation according to the needs of their sectors
 - Though initially non-statutory, the government anticipates introducing a statutory duty requiring regulators to have due regard to the principles in time.
- Delivering new central functions to support regulators to deliver the AI regulatory framework, maximising the benefits of an iterative approach and ensuring that the framework is coherent.