Dilemmas in AI Regulation

An Exposition of the Regulatory Trade-Offs Between Responsibility and Innovation

Markus Trengove[1,2*] and Emre Kazim[2,3]

[1] School of Public Policy, University College London, 29 Tavistock Square, London, WC1H 9QU, UK
[2] Holistic AI, 18 Soho Square, London, W1D 3QL, UK
[3] Department of Computer Science, University College London, Gower St, London, WC1E 6EA, UK
[*] Corresponding author: markus.trengove@holisticai.com

Abstract

In this paper, we map out key strategic and normative dilemmas that regulators must navigate in regulating the development and application of AI. We propose three such dilemmas. The first and fundamental dilemma concerns the tradeoff between Responsibility and Innovation: governments must decide how to balance the imperative for regulatory safety and the imperative to develop and proliferate the use of AI. The first dilemma generates two more specific dilemmas: the choice between Horizontal and Vertical regulation, and the formulation of risk assessment methods. Here, we argue that the latter decision points fundamentally depend on the first: the structure of AI regulation should follow from the relative weighting that regulators give to the competing imperatives of Responsibility and Innovation. The task of this paper is to map out these decision points, and to foreground their normative assumptions and implications.

Keywords:

Regulation; Artificial Intelligence; Compliance; Governance; Legislation; Risk Assessment; EU AI Act; UK Data Strategy; Algorithmic Accountability Act

1. Introduction

National legislators and policy-makers of numerous states are currently attempting to delineate the scope, structure, and substance of regulations relating to Artificial Intelligence (the analysis in this paper applies to AI most broadly construed to include algorithmic processing systems).

The foremost flagship legislative intervention is the EU's proposed AI Act,[1] with the UK following suit with its own vibrant policy, standards, and regulatory debates.[2] The UK government's action

Electronic copy available at: https://ssrn.com/abstract=4072436

has thus far centred around strategic policy documents scoping the government's position on AI. These documents suggest a less interventionist position than the EU. However, the UK's policy strategy still suggests that the key trade-offs in its AI regulation are very much in contention. Lastly, US senators have circulated draft legislation proposing risk-assessments for algorithmic systems.[3] Despite the perception of the US as non-interventionist with regards to technological innovation, the draft legislation promises close scrutiny of the development and risks of algorithmic systems.

Although each jurisdiction has its own context and nuances, it is worth noting that regulation is yet to be codified in each. Accordingly, there remains ample scope for contestation and clarification of the strategic and normative direction of AI policy and regulation in each jurisdiction. The scope for this intervention will widen as the use and relevance of algorithmic systems expand.

The purpose of this paper is not primarily to provide detailed analysis of proposed policy interventions, as we do elsewhere.[4,5] Instead, this paper aims to provide a higher-level conceptual overview of the key strategic and normative decisions that legislators and policymakers will have to make when regulating AI. In this regard, its use of particular national examples should be taken as illustrative, rather than exhaustive.

This paper focuses on three decision points.

- ***Innovation/Responsibility:*** The broader, underlying tradeoff is between Innovation and Responsibility, which we take to describe the ways in which the imperative of facilitating technological discovery and enterprise conflicts with the governmental responsibility to ensure the safe and trustworthy development and use of AI technologies. AI presents numerous moral hazards and it must be developed sustainably and ethically to ensure public safety, as we explore in greater detail elsewhere.[6] However, stringent regulation of AI for the purpose of public safety has the potential to frustrate technological progress. Section 2.1 outlines this conflict, and subsequent sections return to this question in framing further regulatory decisions.

- ***Vertical/Horizontal:*** Second, the paper turns to the decision point concerning the remit of AI regulation: in particular, legislators must decide whether AI regulation will be Vertical or Horizontal. A vertical strategy focuses on sector-specific regulation, often delegated to particular regulatory authorities. Horizontal regulation, by contrast, centralises regulatory norms and procedures by applying generalised sector-agnostic rules and processes to the application of AI. Section 2.2 outlines these approaches and presents their comparative strategic and normative advantages.

- ***Risk Assessment Methods:*** Third, in section 2.3 of this paper, we focus on different models for assessing risks. In deciding the appropriate compliance measures for AI, legislators will have to decide how best to triage risk, so that higher-risk AI is subject to appropriate levels of compliance control, but also so that low-risk AI is not overburdened

Electronic copy available at: https://ssrn.com/abstract=4072436

by regulation. Here, we compare the numerus clausus approach (exemplified by the EU AI Act), and an individualised impact assessment (exemplified by the US Algorithmic Accountability Act). Our intervention here is aimed in part to complement our prior work on the value of risk impact assessments. [7]

Theme	**Conceptual Division**	**Key Takeaway**
Role of Government	Innovation versus Responsibility	Government has conflicting duties to protect the interests of the public and to facilitate innovation
Structure of Regulation	Horizontal versus Vertical	Regulators must decide whether to impose uniform regulation or to defer to industry regulators
Risk Assessment	*Numerus Clausus* versus Impact Assessment	Regulators must decide whether to limit categories of high-risk AI or to assess each case against a proportionality measure

Lastly, in section 3 of this paper, we turn to the question of public interest, which should frame the decisions that legislators and policymakers make in considering the above tradeoffs.

In this paper, we make two critical interventions:

- First, **seemingly technical choices in regulation are ultimately grounded in normative assumptions** about the role of government in interfering in the economy for the public interest;
- Second, we argue that **strategic policy documents have misconstrued the 'public interest' in framing the purpose of AI regulation**.

Here, we preview our forthcoming work on AI and the public interest, to ask how we ought to understand the public interest and how it ought to inform policy decision making with regards to AI.

Theme	Critical Takeaway
Structure of Regulation	- Vertical regulation favours Innovation - Horizontal regulation favours Responsibility

Electronic copy available at: https://ssrn.com/abstract=4072436

Risk Assessment Method	- A *numerus clausus* approach favours Responsibility - A impact assessment approach favours Innovation
Public Interest	- There is no singular public interest: it is composed of competing interests - Innovation is a component part of the public interest that is in tension with rights protection

In brief: in this paper we highlight key trade-offs, expressing them and then offering our commentary. We forward our assessment within the context of, and drawing upon, different regulatory frameworks, ultimately with a view to informing the debate in the UK concerning its own nascent policy and regulatory direction. We conclude by offering a brief critical reflection on these tradeoffs and mapping the space for future intervention on this set of questions.

Our intended audience includes those interested in tech policy–both from the regulatory perspective but also from an industry perspective. Since the effects of tech regulation and innovation are likely to be societal in scale, our intention is also to address members of the public who do not work in tech or policy, but who have a vested interest in the development of the tech industry.

2. Regulatory Dilemmas and Trade-Offs

In this section we offer surveys of the decision points that we take to be the most pressing conceptual questions that should frame legislative and policy interventions in the development of AI. This list is not meant to be exhaustive of the considerations that should inform law or policy. Rather, it is intended as a high-level primer to the fundamental strategic and normative decisions that governments must make in forming their regulatory approach to AI. This section proceed as follows:

- Subsection 2.1 outlines the Innovation and Responsibility Dilemma;
- Subsection 2.2 outlines the Horizontal and Vertical Regulatory Dilemma;
- Subsection 2.3 outlines the Risk Assessment Dilemma;
- Subsection 2.4 provides an analysis of the interconnections between these different regulatory decision points.

2.1 Dilemma 1: Innovation and Responsibility

First, the regulation of AI is subject to a similar tension that underlies most government intervention in the market: the tension between under- and over-regulation. Although this does not hold uniformly, the received wisdom is that under-regulation exposes society to the risks of

Electronic copy available at: https://ssrn.com/abstract=4072436

externalities–most saliently, social, political, and environmental costs–whereas over-regulation might arrest a competitive marketplace, innovation, and economic growth.

The regulation of AI concerns the following: compliance measures that developers are obligated to undertake (e.g. legally-mandated reporting to regulatory bodies concerning the use of personal data); the applications of AI that are legally prohibited (e.g. bans on the use of AI for psychological manipulation); the legal restrictions meant to guide the development of AI technology (e.g. the creation of AI sandboxes meant to facilitate innovation, testing, and validation); and government support for the development of AI (e.g. funding university research and exempting research institutions from stringent regulation) (see our prior work in [5]).

The regulation of AI across these subfields can be more or less stringent. A more responsibility-oriented approach would prohibit a greater number of AI applications, impose greater compliance measures on a greater number of AI applications, impose more stringent legal restrictions on the development of AI, and limit government support for research and development. More stringent measures would likely slow the development of AI and hinder its commercial and governmental applications, but they would also allow for a 'better safe than sorry' approach to AI by taking protective measures to its risks until they are better understood. This approach is consistent with the 'precautionary principle', which recommends prevention and mitigation of potential risks where their likelihood of occurrence is unknown or unverified.

By contrast, a more innovation-oriented approach would allow a broader range of AI applications, limit the number of regulatory compliance measures imposed on permissible AI applications, impose less prohibitive and more encouraging conditions for the development of AI, and provide greater support for research. This approach would facilitate greater innovation in the field by proliferating its use and development without government intervention, but it would also offer less protection against the proliferation of risk occasioned by AI.

This tension is front-of-mind to both the EU and UK regulatory strategies. The EU's approach–exemplified by the EU AI Act and the EU High Level Expert Group on AI reports– is self-characterised by creating 'ecosystems of excellence' and 'ecosystems of trust'. [8] Similarly, the UK, in its National Data Strategy, aims both to 'spur innovation' and 'earn trust'.[2] What is apparent in both these formulations is the competing imperatives that governments recognise: in its capacity as the steward of the national economy, it is incumbent upon the government to stimulate economic growth and innovation; however, the government is similarly charged with respecting and protecting the rights and well-being of its individual citizens.

Neither a fully responsibility-oriented nor a fully innovation-oriented approach would be satisfactory. Veering too much in either direction runs the risk either of economic stagnation or rights-violations. Regulators ideally want to achieve a *via media* that optimises the general interest by balancing growth with responsibility, creating a political economy in which the rising tide of innovation lifts all ships, without exposing ordinary citizens, communities, and consumers to a host of new and unnecessary threats.

Electronic copy available at: https://ssrn.com/abstract=4072436

Framing this in terms of tradeoff, of course, assumes that the two goals are mutually exclusive. There are cases in which state regulation of the market can encourage innovation, by providing breaking up monopolies, assuaging investors of the stability of the market, building a trusting consumer-base, and gaining access to other markets through regulatory alignment. However, for the purposes of this introduction, let us assume that there is a conflict that must be navigated here.

2.2. Dilemma 2: Horizontal and Vertical Regulation

The proliferation of AI is likely to reach almost every reach of the economy, spanning and altering numerous existing sectors. This section discusses how governments should undertake the regulatory division of labour: must AI be regulated by the central government and subject to uniform rules? Or must the authority to regulate applications of AI be limited to industry regulatory bodies imposing sector-specific rules? The former approach of enacting regulations for all applications of AI is what is referred to as 'horizontal' regulation. The latter approach of deferring to industry regulators is known as 'vertical' regulation. This section provides an overview of the comparative advantages of both approaches.

Regulation can be varying degrees of vertical or horizontal with respect to the breadth of its scope and the empowered authority issuing regulations. Horizontal regulation–typically a function of central governments–would issue uniform regulations that apply to AI across industries. Vertical regulation, by contrast, crafts regulation that is sector-specific, or delegates regulative authority to the supervision of an industry body.

There are economic and political advantages to both approaches. The horizontal model is uniform and stable, whereas the vertical model is flexible. The horizontal model enables the central government to guarantee that all sectors meet the standards that are in the public and national interest. It prevents enterprises from 'shopping around' between sectors, ensuring that all players conform to the same standards. Horizontal regulation also gives citizens a stable set of assurances and rights in their interactions with AI–i.e. they know what to expect and will be able to tell when they are being wronged.

By contrast, the vertical model of regulation allows industry bodies to make provision for differences across sectors and to optimise regulation to meet the specific needs of the industry. Vertical regulation has the benefit of specificity, which provides greater certainty to members of the industry where horizontal regulation will necessarily be vague in order to be broad in its application. A vertical approach occasions different coordination problems: there may be a significant amount of overlap between regulations resulting in a multiplicity of responsible agents, overlap and undue burden upon relevant parties; a single party may be required to report to multiple regulators regarding the same thing; the regulators may protect their autonomy and remit, thereby restarting knowledge transfer and and enforcement across sectors.

Electronic copy available at: https://ssrn.com/abstract=4072436

The EU AI Act proposes a horizontal regulatory framework for the European Union, in which AI across sectors are subject to the same risk-assessment criteria and legal requirements (discussed in the following section). Per the AI Act, all low-risk algorithms are subject to transparency requirements, high-risk algorithms are subject to more strenuous compliance measures, and then a further category of algorithms are prohibited altogether.[5] Similarly, the proposed US Algorithmic Accountability Act adopts a horizontal approach by directing the Federal Trade Commission (FTC) to mandate impact assessments of AI systems across sectors (subject to the size and reach of the enterprise).

The UK's approach to AI regulation remains more underdeveloped on this question: whilst the UK has intimated that it favours a sectoral approach, there are indications that possibly point in the direction of horizontal or hybrid regulation. With the move away from the EU, the UK has signalled, in the digital realm, that there is an appetite to propose an alternative approach to digital regulation. This has manifested itself through the activity of individual regulators publishing and signalling in the space of AI. Although there are signals of an alternative UK approach (one that more closely approximates a horizontal approach), we read the current stand as oriented towards a sectoral stance.[9]

The UK Data Strategy,by contrast, is more ambiguous on this point: the Strategy indicates that it will devolve some regulatory powers to existing industry regulators, but suggests too that industry regulation may be underpinned by cross-sector principles that standardise the UK's innovation landscape.[2] This points to the possibility of a hybrid system in which there is vertical regulation supplemented by a minimum core of horizontal regulations that ensure a coherent set of rights, duties, responsibilities, and assurances.

The UK's National AI Strategy[10] can be seen in terms of a (potential) horizontal approach to regulation.[6] We read this into the text because of the sections within the strategy that relate to AI governance. In fact one of the three pillars of the text is 'Governing AI effectively', in which two imperatives are particularly present: stimulating innovation and enterprise, and designing a standards and regulatory regime that reflects this innovation agenda. Within this framework, drawing on our prior work,[6] there are a number of themes to note:

- Removing some existing regulatory burdens where there is evidence they are creating unnecessary barriers to innovation.
- Retaining the existing sector-based approach, ensuring that individual regulators are empowered to work flexibly within their own remits to ensure AI delivers the right outcomes.
- Introducing additional cross-sector principles or rules, specific to AI, to supplement the role of individual regulators to enable more consistency across existing regimes.

Note that these themes read as an *endorsement of the sectoral approach*. In fact, we have read the ecosystem as signalling that the UK will take a sector-specific approach–as such the inclusion of the cross-sector regulatory policy is to be read as an important shift.

That said, our reading also suggests that the UK can position itself for horizontal regulation. To be clear, the language of 'the UK positioning itself' is not a claim that there is a coherent and

Electronic copy available at: https://ssrn.com/abstract=4072436

singular intentional force driving the agenda. Indeed, naturally, there are competing interests and forces that motivate their respective agendas, seeking to realise aims that conform to what they would envision as appropriate and adequate AI governance (whether that be in the form of self-regulation, representation, regulation (horizontal and vertical), etc.). Below we will tease out a narrative that can be thought of as implicitly (if not explicitly) leaving open the possibility of a horizontal form of UK regulation.

Regarding removing existing regulatory burdens, the reading can quite readily be one of moving towards a more principled, national level, regarding regulation. Regarding retaining the sector-based approach, a reading can be made that appeals to the imperative that the regulators should be flexible i.e. the sector regulator burdens may be in setting standards and monitoring rather than in codification of a specific regulatory requirement - something that would readily align with a horizontal approach. As such, an interesting argument can be given that the sectorial/vertical approach can serve a horizontal agenda insofar as the associated regulators of the sectoral/vertical serve as effective policy, standards and monitoring institutions.

Similarly, the UK's AI Assurance Roadmap[11] presents a possible interpretation that suggests a horizontal approach to regulation (see [12]). The text is clear insofar as the aim is to realise mechanisms of ensuring trustworthy AI to establish trust and set the conditions for AI innovation, industry and adoption (very much in the AI as an opportunity genre). In the section that directly addresses 'The Role of Regulators', instruction is given regarding the need for the regulators to provide guidance and expertise, reflecting the nuances of the respective sectors. In closing the section the text reads 'as the use of AI systems becomes more widespread across sectors, clear regulatory scope will need to be established between regulators with similar and overlapping mandates. Moreover, regulators will need to decide in what contexts assurance be directly sought by regulators (regulatory inspection) and when assurance should be delegated to assurance providers, or where assumed conformity is appropriate' - here the reading indicates a point of indeterminacy. Although there is clearly an implicit view that the landscape will evolve through sectoral regulation (insofar as the sectoral regulators and assurance bodies will be tasked in their respective domains), it is still an open question how this will translate to the specificity of regulatory regimes with respect to specific sectors (assurance and compliance are not one and the same thing).

Existing sectors are typically already subject to their own regulatory bodies and regulations that understand the needs of their industries and are well-placed to regulate the development of AI within their sectors. However, the effects of AI are of national (and international) import; many applications of AI cut across sectors; and sector-specific is likely to be piece-meal. Therefore, it is important for regulators to adopt a combination of regulatory strategies that balance the need for industry-specific regulation and national regulatory uniformity.

Electronic copy available at: https://ssrn.com/abstract=4072436

2.3. Dilemma 3: Risk Assessment Methods

The regulation of AI will require some risk triage, so that high-risk AI will either be prohibited or be limited by demanding compliance measures, but low-risk AI will not be burdened by the same regulatory constraints. This raises a methodological question for regulators: how does one perform an *ex ante* risk triage to determine the levels of regulatory and compliance stringency to impose upon different algorithms.

First, one might adopt a *numerus clausus* approach, as in the EU AI Act. At the highest level of risk, Title II of the Act prohibits certain use-cases of AI outright. Next, Annex III of the Act enumerates a list of use-cases that are classified as high-risk. These cases are subject to more stringent regulatory scrutiny, whereas low-risk cases must simply make transparency declarations to users–thus protecting against the most severe harms whilst allowing low-risk developments to flourish without regulatory burden. These use-cases include law enforcement (recidivism assessments; polygraphs; evidence checking).

There are several noteworthy features of this approach. The first is the apparent *raison d'etre* of the list: underlying the use-cases seems to be a concern with the potential violations of civil and basic human rights. The second is that it is, by its nature, limited: the Act makes provision the possibility of adding to the list–presumably as our experience of AI and its possible detriment grows. For this reason, the list is also necessarily retrospective: the use-case list is based on case-by-case experience, and the list enumeration is therefore in some sense an *ex post facto* ameliorative measure. The advantage of such a closed list is that it provides clear guidance to industry by specifying the circumstances in which stringent compliance applies. However, the disadvantage is that it risks lagging behind different risks by only taking preventative measures as threats arise or become actual.

A second strategy is to conduct algorithmic impact assessments (AIA)–as exemplified by the US's proposed Algorithmic Accountability Act. The purpose of an AIA is to survey a number of features of an algorithm, its structure, its governance, and its potential impact, in order to assess the degree of compliance and oversight required. The AIA mandated by the Act is, in essence, a proportionality measure that scores algorithms across a number of metrics, including the following:

- The benefits and need for an AI-based decision-making system, as compared to the previous system;
- The purpose of the AI-based system;
- The risks of harm to consumers associated both with the AI-based system and the preceding system.
- Any differential impacts of the system on the basis of race, gender, disability, and other protected characteristics.
- The transparency and explainability of the AI system and consumers' ability to 'contest, correct, or appeal' the system's decisions.

The AIA also accounts for mitigation factors, including the following:

Electronic copy available at: https://ssrn.com/abstract=4072436

- Consultation with external stakeholders affected by the AI-based decision-making system;
- Internal protocols for the elimination or mitigation of negative impact on consumers;
- Staff training on risk mitigation and protocols.

The AIA includes both ex ante and ongoing phases: enterprises must submit reports for any new applications of AI before deployment, and must submit ongoing monitoring reports on the impact of the system.

Although the Act mandates the FTC to issue guidelines for features of the AIA (including establishing a standardised format for reporting, issuing guidelines for consultation, and issuing guidelines for prioritising systems to be assessed), it seems from the text as though the AIA will be multiply realisable. In other words, enterprises will have a number of possible means of satisfying the AIA test, leaving them with greater liberty.

An Innovation advantage of such an IA measure is that, unlike the *numerus clausus* approach, it is sufficiently flexible to accommodate higher-risk applications of AI as long as they are sufficiently off-set by mitigation strategies. Unlike the *numerus clausus* approach, though, it does not provide clear guidance on the level of acceptable risk, nor is it sufficiently well-defined to preclude the possibility of gerrymandered assessments. If such a proportionality measure is to be effective in mitigating risk, it will have to find a way to quantify relevant risk thresholds that demarcate how much risk warrants what compliance measures.

There are similar suggestions for such impact assessments in the UK's strategic policy documents. The Assurance Roadmap, [13] for instance, leaves it at the discretion of regulators to decide when to assume compliance, when to delegate compliance assurance, and when to seek direct compliance assurance. Reading into this, we can infer that the regulators here will be tasked with making a judgement in which they weigh up the risks of a particular algorithmic application, as well as the mitigation measures in place, to determine the level of scrutiny that will be necessary to manage risk in such a way as to balance safety and innovation.

Governments and regulators need to find ways to assess risks for the purposes of regulation. The form that this assessment takes must be adapted to the normative and strategic imperatives of AI regulation: it must balance the need for clear guidance on the level of acceptable risk with the need for flexible risk management; it must be able to capture the most pressing risks without burdening low-risk developments with onerous regulatory burdens.

3. Public Interest: Returning to the Innovation/Responsibility Tradeoff

In this section, we make two contentions: the first is that the latter two dilemmas (the Horizontal/Vertical choice, and the choice of Risk Triage) are grounded in the

Electronic copy available at: https://ssrn.com/abstract=4072436

Responsibility/Innovation dilemma; the second is that the Responsibility/Innovation dilemma requires a trade-off between competing imperatives within the public interest.

3.1. Innovation and Responsibility Assumptions in Regulatory Choices

Consider the choices between Horizontal and Vertical regulation, and Risk Triage methods. Our concern is that these regulatory choices come to be regarded as technical decisions that should be aimed towards maximising efficiency in regulation. On the contrary, we argue that these choices are fundamentally political: they must be underpinned by a normative view about the role of state intervention in the (digital) economy.

In the choice between Horizontal and Vertical regulation, the Horizontal option favours government intervention, whereas the Vertical option favours non-intervention. Horizontal intervention gives the central government the power to reduce or eliminate risks across the board by issuing universal compliance measures for the development of AI. As we discussed in section 2.2, this comes at the cost of innovation, because it imposes inflexible standards across industries. In this sense, the Horizontal option is more amenable to a Responsibility-oriented approach to AI governance, because it is adept at reducing risk, even if it frustrates innovation. By contrast, the Vertical option is more amenable to an Innovation-oriented approach, because it allows industry regulators to set more adaptable standards that suit the needs of the industry–albeit at the cost of uniform safety.

Similarly, the choice of Risk Assessment method has significant implications for Responsibility and Innovation. First, it is evident that a Responsibility-based approach demands a lower risk threshold, whereas an Innovation-based approach demands a higher risk threshold.If, as the *numerus clausus* approach suggests, regulation prohibits or limits the development of AI in entire categories of use, this will clearly reduce risk, but it will also frustrate innovation by leaving entire fields underdeveloped. By contrast, a proportionality approach leaves open the possibility of innovation in any direction, granted that the benefit of innovation is sufficient to justify taking on some risk. (Of course, in principle, both the *numerus clausus* and proportionality approaches can be more or less internally stringent. A *numerus clausus* approach could, for instance, limit the application of AI in few or no fields. Similarly, a proportionality approach could set its risk threshold so low that few applications will exceed the threshold it sets (in this sense, a proportionality approach depends on regulation's priors; i.e. the comparative importance regulators place on risk and innovation before conducting a proportionality assessment). However, we operate on the assumption here that regulators would use both methods to capture approximately the same set of use-cases. Our claim here is therefore a ceteris paribus assessment.)

In devising regulatory mechanisms, legislators must remain aware of how seemingly technical questions are in fact normative questions about how regulators ought to intervene for the

Electronic copy available at: https://ssrn.com/abstract=4072436

purposes of public interest. This speaks to the broader question of public interest: it is the government's imperative to serve the public interest, which entails cultivating economic innovation and securing a fair rights-respecting society for all.

3.2. Problematizing the Innovation/Responsibility Tradeoff

As we discussed in section 2.1, this tradeoff between Responsibility and Innovation is front-of-mind in the strategic policy documents that different legislators have issues. However, there are some conceptual moves implicit in their framing of the issue that we would like to clarify.

First, policy documents like the UK's National Data Strategy seem to frame Responsibility and Innovation as complementary goals; the Strategy claims that: '[t]he UK's data regime will support vibrant competition and innovation, building trust and maintaining high data protection standards without creating unnecessary barriers to data use'.[2] It should be evident from our foregoing discussion that these are not straightforwardly complementary goals: legislators and regulators will need to make choices that tradeoff risk-averse Responsibility and risk-prone Innovation.

However, there is a second error we wish to avoid: the assumption that 'public interest' is a monolithic category that can be neatly defined in terms of human rights. The EU approach, arguably, errs on the side of rights-protection. This 'human-centric' or 'rights-centric' approach must not be confused with a public interest approach: protecting individuals' rights is only one component of public interest.

Public interest is constituted by an internal tradeoff between the interests of competing groups. In regulating AI, public interest is composed of the interests of disparate groups including tech users, Big Tech, Small to Medium Enterprises, the various branches of government, universities and research institutions, non-governmental organisations, and non-tech enterprises. It should be evident from this list that there will be no uniform set of interests that will capture these groups.

Moreover, although policy documents are right to identify rights-protection as a vital element of public interest, it is also crucial to the interests of the public that the UK harness its economic opportunities in a productive and equitable manner that presents opportunities for enterprising members of the UK, creates employment opportunities, and raises the welfare of all members of society. In this sense, the EU risks adopting a lopsided precautionary principle: one that limits risk when they concern rights, but neglects the risks of costly opportunity costs in frustrating welfare-improving innovation.

Our intention here is not to provide a definitive account of the public interest–although we map the concept in greater detail elsewhere. Our intention is rather to problematize the reliance on the concept that has pervaded policy documents.

Electronic copy available at: https://ssrn.com/abstract=4072436

4. Conclusion

In this paper, we have mapped three key strategic dilemmas that will shape the structure of artificial intelligence regulation: the dilemma between Responsibility and Innovation; the choice between Horizontal and Vertical regulation; and the choice of Risk Assessment Method. Whilst the latter choices seem technical, we argue that they are underpinned by normative assumptions about the role of government in regulating innovation. Moreover, insofar as strategic policy documents frame the role of government as facilitating the complementary imperatives of innovation and responsibility, we argue that this is fundamentally misconstrued: not only is public interest composed of competing interests, it also includes innovation. Therefore, the description of the public interest in policy documents ignores the complexity and competition inherent in the concept–one that we explore further in our paper that complements our intervention here.

Declaration of interest
The authors declare no competing interests.

Author contributions
M.T. and E.K. were responsible for conceptualisation and writing (original draft & review and editing).

5. References

1. European Commission. (2021). Proposal for a Regulation Laying down Harmonised Rules on Artificial Intelligence. https://digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence.\uc0\u160. (Accessed 8 June, 2021)

2. Department for Digital, Culture, Media & Sport. (2021). UK National Data Strategy. https://www.gov.uk/government/publications/uk-national-data-strategy. (Accessed 7 December, 2021).

3. Clarke, Y. D. (2022). Algorithmic Accountability Act of 2022. https://www.congress.gov/bill/117th-congress/house-bill/6580/text. (Accessed 31 March, 2022).

4. Kazim, E., Kerrigan, C., and Koshiyama, A. (2021). EU proposed AI legal framework. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3846898

5. Kazim, E., Almeida, D., Kingsman, N., Kerrigan, C., Koshiyama, A., Lomas, E., and Hilliard, A. (2021). Innovation and opportunity: review of the UK\uc0\u8217{}s national AI strategy. Discov Artif Intell *1*, 14. https://doi.org/10.1007/s44163-021-00014-0

6. Koshiyama, A., Kazim, E., Treleaven, P., Rai, P., Szpruch, L., Pavey, G., ... and Lomas, E.

Electronic copy available at: https://ssrn.com/abstract=4072436

(2021).Towards algorithm auditing: A survey on managing legal, ethical and technological risks of AI, ML and associated algorithms. {\i{}SSRN Electronic Journal}. https://doi.org/10.2139/ssrn.3778998

7.
Kazim, E., and Koshiyama, A. (2021). The interrelation between data and AI ethics in the context of impact assessments. AI Ethics {\i{}1, }219-225. https://doi.org/10.1007/s43681-020-00029-w

8. High-Level Expert Group on AI. (2019). Policy and Investment Recommendations for Trustworthy Artificial Intelligence. https://digital-strategy.ec.europa.eu/en/library/policy-and-investment-recommendations-trustworthy-artificial-intelligence. (Accessed 31 March, 2022).

9. Institute for the Future of Work. (2021). Algorithmic Impact Assessments. https://www.ifow.org/publications/policy-briefing-building-a-systematic-framework-of-accountability-for-algorithmic-decision-making. (Accessed 31 March, 2022).

10. Office for Artificial Intelligence. (2021). National AI Strategy. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1020402/National_AI_Strategy_-_PDF_version.pdf. (Accessed 31 March, 2022).

11. Centre for Data Ethics and Innovation. (2021). The roadmap to an effective AI assurance ecosystem. https://www.gov.uk/government/publications/the-roadmap-to-an-effective-ai-assurance-ecosystem. (Accessed 31 March, 2022).

12. Barrance, E, Kazim, E. et al., Review of the CDEI's Extended Roadmap to an Effective AI Assurance Ecosystem.*Forthcoming*.

13. Ahamat, G., Chang, M., Thomas, C. (2021). The need for effective AI assurance. Centre for Data Ethics and Innovation Blog. https://cdei.blog.gov.uk/2021/04/15/the-need-for-effective-ai-assurance/. (Accessed 13 October, 2021).

Electronic copy available at: https://ssrn.com/abstract=4072436