

The future of AI regulation in Europe

The direction of travel is clear, but questions remain unanswered

In February, alongside its [vision for Europe's digital future](#) and its [strategy for data](#), the European Commission released a [White Paper](#) on artificial intelligence. In putting forward the EC's proposals for regulating AI, it answers some important regulatory questions – but on closer inspection, many of the finer details are left unexplored. Here, we summarise what the White Paper tells us about the economic aspects of the EC's plan, and discuss the factors that regulators will need to confront next.

The White Paper: an overview

The EC's White Paper on AI focuses on the creation of two “building blocks” for the future: an “ecosystem of excellence” and an “ecosystem of trust”.

The ecosystem of excellence concentrates on *policy*, identifying a set of actions to develop AI investment. These include more coordination between member states, the creation of centres of excellence to promote R&D, partnerships between the public and private sectors, and more. The paper is also supported by a series of policy objectives presented in the [European data strategy document](#).

The ecosystem of trust, on the other hand, focuses on *regulation*. A regulatory framework, the EC argues, will be key to promoting the trustworthiness of artificial intelligence – something which is “a prerequisite for its uptake”.¹ It is the EC's proposals on regulation that we'll focus on in this article.

Regulating AI: the proposed framework

Unlike an earlier leaked version,² the published White Paper sets out a neatly ordered framework for regulation, providing clear and consistent answers to the basic questions asked of any regulatory framework:

- **Why** regulate?
- **Who** will be regulated?
- **How** will they be regulated, and **by whom**?

Why?

In regulating artificial intelligence, the EC states that it is seeking to address two major ‘market failures’ in the AI industry:

- **Asymmetry of information:** The complexity of artificial intelligence makes it difficult to scrutinise the development and application of AI systems. This lack of transparency hampers the uptake of AI by businesses, consumers and public authorities.
- **Negative externalities:** many commentators have identified various risks in relation to the safety and the potential harm to consumers of AI applications. Also, these risks can potentially hamper the uptake of this technology and, hence, its success.

Reflecting on these market failures, the paper states that “the use of AI brings both opportunities and risks.” While acknowledging that “AI can help protect citizens’ security”, the EC are aware that “citizens

¹ Page 1.

² Available at: <https://www.euractiv.com/section/digital/news/leak-commission-considers-facial-recognition-ban-in-ai-white-paper/1419563/>

also worry that AI can have unintended effects or even be used for malicious purposes.” Indeed, “lack of trust”, according to the White Paper, is one of the main factors “holding back a broader uptake of AI.”³

What?

With the proportionality principle of EU policy-making in mind, the Commission proposes to divide AI applications into two categories: ‘high risk’ and ‘low risk’. Regulation will **only apply to high-risk applications**.

An AI application will be deemed ‘high risk’ if two (cumulative) criteria are met:

1. It is employed *in a sector* where significant risks might occur (e.g. healthcare, transport, energy and parts of the public sector).
2. It is *used in such a manner* that significant risks, which cannot reasonably be avoided by individuals of firms, might arise (e.g. the risk of injury, death, or significant material or immaterial damage).⁴

In our opinion, these criteria are sensible, and strike the right balance between being too imprecise and too prescriptive. But future iterations of this proposal would benefit from a more explicit reference to economic theory – and also to the importance of quantifying the risks in point 2, and weighing them up against the costs of implementing, and complying with, regulation.

Who?

In short, regulation will apply to companies providing high-risk AI applications in the EU. But the White Paper raises some interesting further details.

Companies, for example, who are based outside Europe will be subject to regulation if they provide applications within EU boundaries. This raises the possibility of the EC’s AI regulation taking on a pan-European form – some companies may adopt it as a standard to apply in other jurisdictions outside the EU (as is already happening with GDPR).

The paper also states that regulation will be applied to those who are “best placed to address any potential risks.” It outlines, for example, that: “While the developers of AI may be best placed to address risks arising from the development phase, their ability to control risks during the use phase may be more limited”.⁵ So, while no reference is made to cost–benefit analysis, or to economic impact assessment, it’s clear that these tools will be central in deciding which segments of the vast and often fragmented AI industry will be affected by future regulation.

How?

The regulations (called “requirements” in the White Paper) the Commission is proposing to apply are as follows:⁶

- *Training data*: When using data to train algorithms, companies must ensure that the EU’s values are respected. Data must be sufficiently broad, covering all potentially dangerous scenarios – some AI literature, for example, identifies data as the main source of bias in algorithms.⁷
- *Record keeping*: Proper records must be kept of data used to train high-risk AI systems, including data sets (if necessary) and documentation on programming and training methodologies. This ‘audit’ requirement is similar to that being considered by competition regulators in the digital advertising sector (e.g. [the ACCC in Australia](#)).
- *Proactive provision of information*: Citizens should be informed, for instance, when they are interacting with an AI system rather than with a human being (unless this is already obvious).

³ Page 9.

⁴ Page 17. There are exceptions: recruitment processes and biometric identification (and other intrusive surveillance technologies) will be considered high risk.

⁵ Page 22.

⁶ See pages 18–22.

⁷ See: Bo Cowgill and Fabrizio Dell’Acqua, Columbia University. *Biased Programmers? Or Biased Data? A Field Experiment about Algorithmic Bias*. Presentation at NBER Conference on Economics of Artificial Intelligence, Toronto, 26–27 September 2019.

- *Robustness and accuracy*: This requirement states that “all reasonable measures should be taken to minimise the risk of harm”.⁸ While all the paper’s regulatory principles are by necessity quite high level, this requirement could benefit from slightly more detail on, at least, what the regulator means by “robustness and accuracy”, in terms that can be recognised by industry players.
- *Human oversight*: This, argues the White Paper, helps to “ensure that an AI system does not undermine human autonomy”.⁹ This might mean, for example, that the output of an AI system does not become effective unless it has been previously reviewed and validated by a human being. Human oversight can be *ex ante* or *ex post*, and will have different levels of conditionality depending on the risk posed by the application.
- *Biometric identification*: Current legislation indicates that AI can only be used for remote biometric identification when this is duly justified and subject to adequate safeguards. Nonetheless, the Commission intends to launch “a broad European debate” on the subject.¹⁰

As to how these requirements will be implemented, the White Paper states that **conformity assessments** will be carried out. In other words, when a company or public institution seeks to implement an AI-based product or service, they’ll need to complete an assessment with the relevant authority.^{11,12}

By whom?

The EC’s paper expresses a clear preference for regulation by *existing* institutions – as opposed to the creation of a new regulator. The Commission repeatedly highlights that existing testing, inspection and certification infrastructures are a model that could be applied directly to AI regulation.

Interestingly, the earlier leaked version of the White Paper appeared more in favour of creating new, AI-specific regulatory bodies. “It would be appropriate”, this version stated, “to appoint authorities responsible for monitoring the overall application and enforcement”.¹³

Where next? The unanswered questions

The White Paper, as we’ve seen above, provides us with an adequate overall picture of the proposed regulation of AI. But when we dig a little deeper, there remain a number of unanswered questions and unexplored details.

Is there scope for more flexible regulation?

The EC seems to propose a kind of ‘command and control’ system of regulation – in other words, the regulators will decide how companies comply with requirements. This is as opposed to ‘economic incentive’ regulation, where a system of monetary penalties and rewards are put in place to incentivise firms to behave in a certain way.

‘Command and control’ does appear to be more fit for purpose in the case of high-risk AI systems. But given the uncertain and rapidly changing nature of AI, is there any scope for more flexible regulatory approaches, such as self-regulation or regulatory sandboxes? If so, in which areas?

What are the other potential dangers of AI?

Rights and safety are not the only areas where AI systems can affect the wellbeing of citizens. Some of the concerns about AI, it is true, may be dealt with by other institutions – competition authorities, for example, may address potential consumer abuse by discriminatory pricing. In our opinion, however, the White Paper could have conducted a deeper analysis of the potential issues that may require regulation – even if they will ultimately be tackled other organisations.

Where might some of these further risks be found? They relate to:

⁸ Page 20.

⁹ Page 21.

¹⁰ Page 22.

¹¹ This assessment may take place on a recurring basis over the lifetime of the AI system, as it “evolve[s] and learn[s] from experience” (page 23).

¹² For low-risk AI systems, a quality label indicating that the system is trustworthy would be developed. Participation in the label system would be voluntary, but if a developer used the label, the requirements would be binding (page 24).

¹³ Page 18 of the leaked White Paper.

- Horizontal and vertical coordination between firms
- State aid¹⁴
- Taxation of activities and innovations related to AI (where the generation of added value is even more vague than elsewhere in the digital world)
- Labour markets¹⁵
- Public procurement (in relation to the ambition of using public sector organisations as early adopters of AI technologies).

Should there be a specialised regulator?

The White Paper would have benefited from more analysis on the advantages and disadvantages of a specialised regulator. This was, perhaps, outside the paper's remit, but it's an important question, and one that remains unanswered. Given that some regulators are horizontal (i.e. they cover multiple industries) and others vertical (cover specific industries), there is a risk that some AI requirements will 'fall through the cracks', or that regulation may be duplicated.

Recent research indicates that company managers believe regulation will limit AI adoption.¹⁶ Putting the correct framework in place is therefore key – and this includes the institutions in charge of governance.

Similarly, while it's true that the EU is at the forefront of product safety regulation, further thinking is necessary on whether existing processes can work in the much more complex and fast-paced world of AI.

What about the details?

Many of the proposals in the White Paper (and in the data strategy) are expressed in very broad terms. On the one hand, this is understandable at this early stage of the policy-making process, with an extensive consultation exercise likely to follow. But on the other hand, some proposals would have benefited from a greater level of detail. After all, the impact of this regulatory package will be determined by the precise features of its design and implementation. For example:

- Careful consideration will need to be given to the costs that regulation will impose, in order to avoid unduly affecting small and medium-sized enterprises. Human oversight requirements, for example, are likely to be particularly burdensome for smaller companies.
- Most of the regulatory requirements – in particular, changes to companies' legal liabilities – will need to be coordinated globally, to minimise the risk of 'scaring away' potential investors.
- During the consultation phase, it will be vital to avoid under-regulating some sectors and over-regulating others. Given the complexity of the industry, policymakers will need to rely on external experts to help.
- Regulators will also need to strike the right balance between data privacy considerations, distributional impacts, and the breadth and depth of transparency, safety and audit requirements for AI systems. Thorough, evidence-based economic analysis will be fundamental to this, in order to identify the 'regulatory sweet spot' (see, for example, [Frontier's analysis](#) for Citizens Advice, of the costs and benefits of personalised pricing).

How will regulation link to other policies?

Alongside the White Paper and data strategy, the EC has published a [more high-level document](#) outlining their vision of digital markets and global competitiveness. It highlights a series of policies and initiatives due to take shape in the next few months. The White Paper does not, however, go into detail about how these policies will interlink with the regulatory framework.

The policies in question include:

¹⁴ See pages 6, 14 and 16 of the [data strategy communication](#).

¹⁵ See, for example, Frontier Economics' [report](#) on what economic evidence says about AI and work.

¹⁶ Mariano-Florentino Cuéllar, Benjamin Larsen, Yong Suk Lee and Michael Webb. *Impact of Artificial Intelligence Regulation on Artificial Intelligence Adoption and Innovation*. Presentation at NBER Conference on Economics of Artificial Intelligence, Toronto, 26–27 September 2019.

- The creation of EU data spaces in nine strategic markets for European firms, the launch of an EU strategy for standardisation, and the development of the Digital Education Action Plan.
- The announcement of the Commission's industrial strategy, due in early March, which is expected to contain a series of initiatives promoting the development of 'AI champions' at a European level.
- An ongoing evaluation of the fitness of EU competition rules for the digital age (2020–2023), and the upcoming launch this year of a sector inquiry into digital markets. These reviews will lead to recommendations that will likely have an impact on companies in the data and AI sectors.

Conclusion

The EC's White Paper, accompanied by its data strategy and vision for Europe's digital future, provides us with a clearer understanding of the EU's direction of travel on the regulation of artificial intelligence. It also outlines the basic shape that the regulatory framework will take.

But, thanks to the White Paper's high-level nature, important questions remain unanswered. As such, a concerted and rigorous effort from institutions and market participants will be required, if the EU is to become a role model for proportionate, effective and evidence-based regulation in innovative markets.

Federico Bruni

 +32 2 896 4875
 +32 470 70 25 09
 federico.bruni@frontier-economics.com

José María Rodríguez

 +34 91 343 2355
 +34 664 47 1016
 jose.maria.rodriguez@frontier-economics.com