Axon Advisory Research

The European Commission's proposed AI Act: a global regulatory paradigm?

The European Commission recently published its draft "AI Act", the first example of proposed legislation for AI globally. Following previous trends in technology law, this is likely to have a ripple effect on other jurisdictions, at least by offering a model for similar initiatives. In this paper, we take a deep dive into the implications of the AI Act for regulators inside and outside the EU.

Axon Partners Group

May 2021

www.axonpartnersgroup.com

Authors:

George Metaxas,
Legal & Regulatory Expert

Samuel Tew,
Senior Manager

Panagiotis Trakas,
Associate

1. Introduction

The use of Artificial Intelligence (AI) in business is increasingly prevalent and disruptive. According to a recent study by ONTSI[1], up to 7% of all SMEs and 24% of all large companies in the European Union ("EU") were using some form of AI in 2020. This is a vast number of businesses and will only increase as the technology becomes more accessible and affordable, and competition renders it a necessity. Similar trends should be also expected in developing countries: software travels easily.

AI has become a favourite buzzword for investors in recent years, with VCs in particular increasing its prevalence in their portfolios.

It is no surprise that AI has become a favourite buzzword for investors in recent years, with VCs in particular increasing its prevalence in their portfolios. According to Investment Platform Dealroom, VC investments in Artificial Intelligence have grown substantially across most regions in the past 5 years, with 30% annual growth in both Europe and the MENA region.

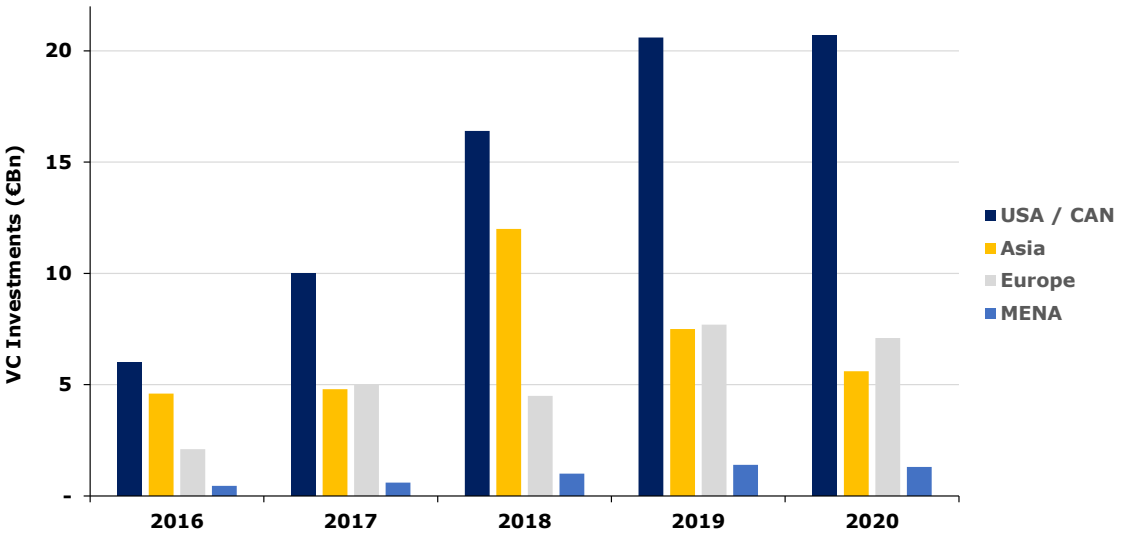

Exhibit 1.1 - VC Investment by region (excl. financial sector)
[Source: Axon based on Dealroom[2]]

However, with great power comes great responsibility: AI's power also carries with it a large propensity for doing harm, especially in large scale applications, as illustrated by the following examples:

[1] ONTSI (Spanish national observatory of Technology and Society), 2021, Indicadores de uso de Inteligencia Artificial en las empresas españolas (red.es), Data taken from Eurostat.

[2] Calculated using Dealroom's online investment platform, 2021, dealroom.co.

- *Gender bias*: Amazon's use of an AI algorithm to hire software engineers in 2015 did not, reportedly, treat both genders similarly as it had been trained with data from recent years in which the number of male candidates had been much higher.[3]
- *Racial bias in the healthcare industry*: In 2019, an AI algorithm was employed to predict which patients would require extra medical treatment. The result was surprising, predicting that white patients would need substantially more medical treatment. This was due to the consideration, as input to the algorithm, of the expenses incurred for the same condition on average, which white patients were presumably more able to afford, as being more affluent on average.[4]
- *Abuse of rights*: Real-time surveillance technologies have been in use in China since 2019 to control and monitor the population by means of more than 200 million cameras. The social control achieved through this real-time surveillance is considered particularly controversial in the West.[5]
- *Situational complexity rules*: An example of an AI technology and a sector that requires a clearer and more adapted regulation is the autonomous vehicle. For example, there is the famous case of a "social dilemma" for these autonomous vehicles, in which their priority is to reduce traffic accidents, but they are presented with an extreme choice between running over pedestrians or crashing the car, thus sacrificing or endangering the lives of the car's occupants or those of others.

AI's power also carries with it a large propensity for doing harm.

Governments and authorities across the world have been avoiding legislating on such issues, suggesting policy or voluntary measures, or non-AI specific rules as a solution instead.

Governments and authorities across the world have been avoiding such issues, suggesting policy or voluntary measures, or non-AI specific rules as a solution. Our research and ongoing coverage of AI initiatives has identified several examples of national strategies,[6] official guidance to

[3] Source: Amazon scraps secret AI recruiting tool that showed bias against women | Reuters

[4] Source: Real-life Examples of Discriminating Artificial Intelligence | by Terence Shin | Towards Data Science

[5] Source. How China Is Using Facial Recognition Technology: NPR

[6] E.g., https://english.msit.go.kr/SYNAP/skin/doc.html?fn=14acc067ebaf2780a558e24993a560f0&rs=/SYNAP/sn3hcv/result/ for South Korea.

state authorities,[7] recommendations,[8] model governance frameworks and principles,[9] proposed Government AI principles[10] etc.

Sooner or later, things were bound to move beyond such policy measures or non-binding recommendations.

Sooner or later, things were bound to move beyond such policy measures or non-binding recommendations. One of the key international rule-makers has now taken the plunge and proposed AI legislation "with teeth": on 21 April 2021, the European Commission (the "Commission") unveiled a draft Regulation laying down AI rules. It refers to it as the Artificial Intelligence Act (the "AI Act").

The world's first AI Act, although still a Commission proposal and bound to be modified in the course of the complex EU legislative process, has sufficient momentum to be adopted in the near future - perhaps by the end of 2022. With the US still scrambling, and so far failing, to adopt federal data protection legislation, the EU will likely benefit from an opportunity to set the legislative agenda in the AI field too – as it has for personal data protection or the ex-ante regulation of electronic communications, to mention two examples in neighbouring areas.

In this article, we provide an overview of the AI Act's provisions and the reasons why its implications will not be limited to the EU territory, but must be considered seriously by AI stakeholders, policy makers and regulators across the globe.

[7] E.g., https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-06.pdf for the USA.

[8] E.g., https://www.priv.gc.ca/en/about-the-opc/what-we-do/consultations/completed-consultations/consultation-ai/reg-fw_202011/ for Canada.

[9] E.g., https://www.pdpc.gov.sg/Help-and-Resources/2020/01/Model-AI-Governance-Framework in Singapore; http://chinainnovationfunding.eu/dt_testimonials/publication-of-the-new-generation-ai-governance-principles-developing-responsible-ai/ in China.

[10] E.G. https://www.smartdubai.ae/initiatives/ai-principles for Dubai; https://www.industry.gov.au/data-and-publications/building-australias-artificial-intelligence-capability/ai-ethics-framework/ai-ethics-principles for Australia.

2. What exactly is AI?

The term "Artificial Intelligence" is nowadays commonly accepted and loosely used in business, government, and industry, even if its exact meaning is only truly understood by the more technical crowd. This may be fine for the inner workings of an AI system, but there is no good reason why the concept of AI should remain engulfed in mystery for the layman, especially if it is associated with rights and obligations that affect the public at large. Laws that bite cannot rely on fuzzy buzzwords.

There is no good reason why the concept of AI should remain engulfed in mystery for the layman, especially if it is associated with rights and obligations that affect the public at large.

A common misconception of AI is that it is somehow "alive". Contrary to "traditional" computing, AI allows machines to imitate, to some extent, human cognitive abilities, and learn from their own "experience", creating an impression of human-level intelligence. But while AI can emulate human tasks such as image recognition or value prediction with impressive efficiency, it is ultimately no more than a combination of computer algorithms based on statistical computing that processes a series of inputs to give a certain type of output.

One of the anthropomorphic concepts involved in AI is that of a "neural network", i.e., a complex set of algorithms that aims to recognize underlying relationships in a set of data, through processes that mimic the architecture and operation of the human brain. Our brain consists of "neurons" that are interconnected to, and can be triggered by, each other. The equivalent (virtual) "neurons" in an AI neural network are used for data input, intermediate processing, and output, as shown in the figure below.

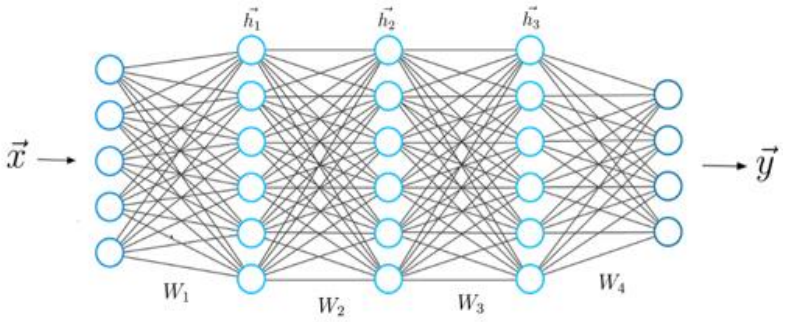

Exhibit 2.1 - Depiction of an AI neural network [Source: Towards Data Science[11]]

[11] Towards Data Science, "What is Deep Learning and How does it work?" (2019). Link: https://towardsdatascience.com/what-is-deep-learning-and-how-does-it-work-2ce44bb692ac

Neurons carry "weights" and (constant) "bias values" which, in combination, determine the neuron's output from a given input. Both weights and bias values are learnable parameters and neurons must be "trained" so that these parameters are adjusted accordingly. In recent years, faster processing – and hence faster training - and big data have led to "deep learning" systems with impressive results, initially in image recognition, and nowadays in a range of other applications.

Therefore, and although the term "artificial intelligence" was coined in 1956, it is only in recent years that it has been matched with sufficient computer power, hardware and data to start delivering on its earlier promises. As such, the potential scale and power of AI use cases can help explain the EU regulators' view that the time is ripe for catching up on the regulatory front.

Although the term "artificial intelligence" was coined in 1956, it is only in recent years that it has been matched with sufficient computer power, hardware and data to start delivering on its promise.

3. Unpacking the AI Act

The AI Act is a long text, with 85 articles and several annexes. These raise three key questions of interest to AI stakeholders and regulators across the world:

- How can AI be legally defined? Is this definition indispensable?
- Which AI systems should be regulated?
- How does the AI Act's risk-based approach work and does it offer a model for other regulators?

3.1. Can AI be legally defined?

The first part of any law or contract is supposed to deal with definitions, but this already poses some serious difficulties in the case of AI.

The Commission claims that "the definition of AI system in the legal framework aims to be as technology neutral and future proof as possible" – but is this really true?

For a start, the definition of "AI system" in the AI Act is general enough to apply to pretty much any computer program: "*software that is developed with one or more of the techniques and approaches listed in Annex I and can, for a given set of human-defined objectives, generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with*". It is therefore the "techniques and approaches" exhaustively listed in the Regulation's Annex I that one needs to look at. At least at first sight, these do not seem to be technology-neutral and future-proof (and arguably not very clear to the layman either). Specifically, the Commission lists the following in Annex 1 of the draft regulation:

The European Commission's definition at this stage seems comprehensive and sufficiently abstract to cover existing AI techniques – but in fact it may be too abstract and therefore overinclusive.

- Machine learning approaches, including supervised, unsupervised and reinforcement learning, using a wide variety of methods including deep learning;
- Logic- and knowledge-based approaches, including knowledge representation, inductive (logic) programming, knowledge bases, inference and deductive engines, (symbolic) reasoning and expert systems;

- Statistical approaches, Bayesian estimation, search and optimization methods.

This list at this stage seems fairly comprehensive and sufficiently abstract to cover existing AI techniques – but in fact it may be *too* abstract and hence overinclusive: for instance, would every computer system involving "statistical approaches" qualify as AI? Obviously not: a simple Excel spreadsheet may also be used for "statistical approaches" and yet nobody would think of qualifying it as AI. This basic example illustrates a first problem with any attempt to provide a legal definition of AI: the individual elements, techniques etc. that AI can ultimately be divided into can be trivial and are shared with a broader spectrum of computer systems and techniques that have nothing to do with AI. It is the sheer ***scale, processing speed and complex interaction*** of these "ingredients" that elevates AI to a new level, even if the dividing line between AI and non-AI computer systems remains fuzzy. Defining AI based on its individual "ingredients" may be as misdirected as trying to define a Da Vinci painting as a "set of paint strokes", putting it on an equal footing with a toddler's worthless smudge (or a contemporary artist's very pricey smudge...).

It is the sheer scale, processing speed and complex interaction of AI's "ingredients" that elevate it to a new level, even if the dividing line between AI and non-AI computer systems remains fuzzy.

One way around this problem may be to avoid an exhaustive legal definition of AI or the "techniques", "approaches" etc. that it is comprised of, and instead focusing on specific use cases to be regulated, based on a risk assessment as the main test. As an example, in one of the proposals submitted to the Canadian government last November as part of its Consultation on Artificial Intelligence, an argument was made that any legislation on AI should not seek to define it as a technology concept but instead focus on the legal implications, such as human rights issues, that come out of the use cases.[12] Encouragingly, as will be discussed below, the Commission's proposal is ultimately also about use cases, even if it has not avoided the temptation of a (probably futile) definition of AI.

For now, the proposed Regulation's definition of AI is not just overinclusive; it can also turn out to be under-inclusive in the future. Currently, AI use cases exist only in the form of "narrow" (or "weak") AI, as programmes designed to do one thing very well (e.g., play chess, identify images, translate text, etc.). This is much narrower than the

[12] Policy Proposals for PIPEDA Reform to Address Artificial Intelligence Report (2020). Link: Policy Proposals for PIPEDA Reform to Address Artificial Intelligence Report - Office of the Privacy Commissioner of Canada

concept of "general" (or "strong") AI, whereby a machine can be presented with any problem or field (e.g., a game, a science, a type of visual input) and self-develop the capabilities to master it without further programming.

Whilst this latter concept is not yet in existence, reaching it is the ultimate goal of the AI industry. It may well be the case, however, that general (or strong) AI one day arrives through the successful implementation of a newly defined form of technology (or "approach"), not covered in the Commission's definition. A concern in that direction can be found, for example, in the Japanese government's *Draft AI R&D Guidelines* (2017), which referred to narrow AI as the basis of their definition but clarified that "*how to define AI in the Guidelines needs to be continuously discussed based on the trends of the technological progress of AI, etc.*"[13]

Striking the right balance between "horizontal" and use-case based rules will be one of the main challenges facing lawmakers on the road to a successful regulation of AI

It is relatively easy to regularly update non-binding guidelines, voluntary codes, etc., but it will be much more cumbersome to adjust an EU Regulation to future technological innovation. Anticipating this challenge, the Commission has proposed that EU Member States empower it to directly update the given list of AI techniques and approaches itself through "delegated acts". It remains to be seen if Member States will accept delegating any such power, even if from a European perspective it would make obvious sense for them to do so.

As a takeaway for any non-EU authorities tempted by the idea of regulating AI, it seems particularly challenging, if not illusory, to define AI in a legally binding, precise and exhaustive way, even if this would normally be a logical first step in any legal text. Instead, the more sensible legal approach may be to focus on specific use cases and the concrete risks these may present, and forget about capturing and policing all of AI.

Admittedly, such an approach would not be risk-free either: breaking up AI rules by different use cases, with no overarching definitions and principles applying horizontally to all types of AI may lead to a fragmented and incoherent regulatory approach, with different regulators and policy makers vying for different pieces of the AI cake, through different legal and policy means.

[13] Japanese Ministry of Internal Affairs and Communication. A tentative translation can be found at: 000507517.pdf (soumu.go.jp)

In conclusion, striking the right balance between "horizontal" and use-case based rules will be one of the main challenges facing lawmakers on the road to a successful regulation of AI.

3.2. Which AI systems will be caught?

All AI systems falling under the Act's broad definition will be caught, with some exceptions.

All AI systems falling under the Act's broad definition will be caught, with some exceptions: the AI Act will not apply to military AI systems or those used for international law enforcement and judicial cooperation with the EU. It will also not apply directly to high-risk AI systems that are safety components of certain products (e.g., those used in civil aviation, rail and marine systems and equipment, and various types of motor vehicles). However, once the AI Act is adopted, the substance of its provisions will be integrated into the existing EU sectoral safety legislation covering these excluded products too, to ensure consistency.

Importantly for non-EU industries and regulators, the AI Act will also apply outside the EU, to:

- Any providers placing AI systems in the EU market, even if they are established outside the EU;
- Any users of AI systems located within the EU; and
- Any providers and users of AI systems located outside the EU, where the output produced by the AI system is used in the EU.

Any AI system run on a non-EU server or cloud, but accessible by a user located in the EU, or producing output that is used in the EU will be also caught by the AI Act.

These are broadly described situations, and their exact implications for the intrinsically borderless AI technology will take time and, possibly, litigation to sort out. At this point, however, it looks as if the AI Act can have significant spill-over effects outside the EU, with implications for any country endeavouring to develop world-class AI systems and products.

As is clear from the above list, any AI system run on a non-EU server or cloud but accessible by a user located in the EU will be caught by the AI Act. Moreover, this will be the case even if the AI system user is *not* located in the EU, as long as "the output produced by the AI system is used in the EU". This is an alarmingly vague condition: how far should one go in defining "output" and what exactly would its "use in the EU" consist in?

Such problems could be heightened by the absence of a one-stop-shop mechanism under the AI Act: contrary to GDPR, there will not be a leading

national authority for AI, at least under the current text. As a Regulation, the AI Act will be directly applicable in all EU Member States, with no need for national implementation measures. However, its interpretation will be left to national courts and/or other national authorities, and can lead to different solutions across the EU. To minimize such diverging interpretations, the AI Act includes provisions for the creation of a European Artificial Intelligence Board (EAIB), which will supervise and facilitate the consistent application of AI legislation and share best practices.

At a minimum, third country AI regulators and policy makers should consider, as a matter of internal due diligence, any risks of a future conflict with the AI Act's provisions, based on their current AI strategies.

In light of the AI Act's extra-territorial implications and Europe's weight in the global economy, other jurisdictions may feel compelled to align their own rules with at least certain aspects of the EU AI rules. At a minimum, third country AI regulators and policy makers should consider, as a matter of internal due diligence, any risks of a future conflict with the AI Act's provisions that could cut short benefits aspired to through their national AI initiatives.

As an obvious precedent for such extraterritorial spill-over, one can mention the EU's data protection legislation, which has led to regulatory adjustments in other countries to ensure the free flow of personal data with the EU or simply to rely on the GDPR as a new data protection "gold standard" worth emulating. Recent examples that come to mind include South Korea's reforms to its data protection laws, which have led to the successful conclusion of "adequacy talks" with the European Commission in April this year; and the Dubai International Financial Centre's (DIFC) EU-inspired Data Protection Law No 5 of 2020 and Data Protection Regulations.

3.3. The AI Act's Risk-Based Approach

The EC's basic approach is to distinguish between three categories of AI systems and practices and to regulate them based on their level of risk.

The EC's basic approach is to distinguish between three categories of AI systems and practices and to regulate them based on the level of risk they present: (i) unacceptable, (ii) high and (iii) low or minimal.

AI practices presenting an unacceptable risk are simply prohibited. Such practices include manipulation through subliminal AI techniques, the exploitation of vulnerable groups, social scoring by public authorities, and real time remote biometrical identification in public spaces for law enforcement purposes (albeit subject to allowed exceptions

for missing children, addressing specific and imminent threats to life, safety or threats of a terrorist attack and for law enforcement against certain specified crimes).

AI systems presenting a high risk are subject to certain mandatory requirements and an ex-ante assessment. This list of high-risk AI systems includes, first of all, AI safety components or complete products covered by certain EU harmonising measures, insofar as they are subject to third-party conformity assessment under EU law before being put on the market. The relevant EU assessment measures (19 in total) listed in an annex to the AI Act cover a wide range of products, not always associated with AI, e.g., toys or elevators - but also radio equipment, medical devices, civil aviation equipment, various types of vehicles, etc. which are more clearly associated with AI.

The AI Act's lists of high-risk AI systems cover a wide spectrum of industries and provide a framework of reference also for non-EU stakeholders in the AI industry.

A separate annex lists further AI systems considered high-risk. Examples include those used for biometric identification; management and operation of civil infrastructure; education and vocational training; employment; certain essential private and public services; law enforcement; etc.

These lists are long and do not need to be replicated here. Suffice to say that they provide a framework of reference also for non-EU producers planning to offer AI systems or products in the future, as well as non-EU policy makers wishing to support the development of such AI systems or products in their country. Given the wide jurisdictional net thrown by the AI Act and the very real prospect of "EU regulatory spill-over", a careful review of the annexes' lists should be part of any AI system or policy developer's due diligence to determine the extent to which the AI Act may affect its business, even outside the EU.

High-risk AI systems will be subject to a detailed list of requirements. In sum, these include:

- a risk management system for the entire lifecycle of each high-risk AI system;
- specified quality criteria for training, validation and testing data sets;
- up-to-date technical documentation to be drawn up already before the AI system is placed on the market;
- record-keeping, including logging capabilities;

- transparency and provision of information to users;
- human oversight interface tools and measures;
- design ensuring accuracy, robustness and cybersecurity, meeting specified requirements.

In addition, the AI Act includes a set of obligations for high-risk AI system providers and users, as well as other parties (such as importers and distributors), largely aimed at ensuring that the above list of requirements is properly adhered to. Further, each EU Member State must designate a notifying authority responsible for the various notifications required under the AI Act.

The AI Act includes unusually detailed provisions on enforcement, with fines that can reach up to €30 million or 6% of the offending company's total worldwide turnover.

Finally, high-risk systems should be subject to a list of post-market monitoring, information sharing, and market surveillance measures. These include incident-reporting and withdrawal or recall obligations.

The AI Act's provisions are more relaxed about AI systems raising **only low or minimal risks.** They essentially require that natural persons should be informed that they are interacting with an AI system, unless this is obvious from the circumstances – with exceptions for legally authorised AI systems used for law enforcement. "Deep fakes" should be also disclosed.

Last but not least, the AI Act includes unusually detailed provisions on enforcement with administrative fines that can reach up to € 30 million or 6% of the offending company's total worldwide turnover, whichever is higher.

4. Conclusion

To recapitulate, the AI Act will likely have repercussions for any future legal regimes applying to AI across the world, even if its territorially scope is nominally limited to the EU. First, the AI Act is the first example of proposed legislation in this area, proposed by one of the global regulatory leaders, and hence likely to offer a model for similar initiatives by other jurisdictions: if a set of AI rules works for 27 EU Member States after going through the gruelling scrutiny of the EU legislative process, it is likely to also offer a credible model to consider outside of the EU too.

The AI Act is the first example of draft legislation in this area, proposed by one of the global regulatory leaders, and is therefore likely to offer a model for similar initiatives by other jurisdictions.

Second, the AI Act will also apply to AI systems developed or operating outside the EU insofar as these have an impact inside the EU, based on broadly scoped criteria. For example:

- AI-based CV-analysis software developed and used by a multinational organisation somewhere in the Middle East or in the Americas could face problems assessing candidates for its EU subsidiaries or assignments if it is not AI Act compliant;
- Developers of AI systems for smart energy grids or smart cities will not be able to sell their systems in the EU unless they adjust them to the AI Act's requirements;
- Non-EU telecoms operators using AI systems in their roaming services could face legal issues handling the data of visiting roamers from the EU.

As shown by such examples, ambitious national AI programmes, systems and solutions outside the EU could turn out to miss their objectives if their output is unusable within the EU or by remote users based in the EU. This is why the principles proposed in the AI Act should be the subject of close monitoring and early consideration both within and outside the EU.

The AI Act is still in a draft stage, likely to undergo substantial modifications going forward and will not be perfect. Some of its principles, such as an attempt to provide a legal definition of AI, may turn out to be particularly challenging. Nevertheless, its general risk-based regulatory approach makes sense and can offer a model for other jurisdictions too, despite any tensions between different policy priorities (e.g., ethical principles, public security concerns and business interests) already reflected in the first reactions to the proposal.

Its general risk-based regulatory approach makes sense and can offer a model for other jurisdictions too, despite any tensions between different policy priorities.

5. About Axon Advisory

Axon, through its Advisory arm, is an international investment and advisory firm offering world-class consulting and corporate finance services to a broad client base in the broad technology sector.

In the last 10 years, Axon has executed +500 projects in +60 countries for major private companies, governments, institutional bodies, and technology companies worldwide.

Axon has been at the forefront of ICT regulation, having advised policymakers and regulators across the globe in a wide range of issues, including the definition of policy and regulatory roadmaps and the update of regimes in the new digital ecosystem.

Analysts Team at Axon Partners Group[14]

[14] The views and opinions expressed in this article are those of the authors and do not necessarily reflect the view of Axon Partners Group.