COMPLIANCE WITH THE REQUIREMENTS (ARTICLE 8) *

Alejandro Huergo Lora
Full Professor of Administrative Law
University of Oviedo

CONTENTS: 1. Some preliminary remarks. 2. Criteria to be taken into account in the assessment of compliance with the requirements (paragraph 1). 3. Compliance with the requirements regarding to products to which requirements of the Union harmonisation legislation listed in Section A of Annex I apply (paragraph 2).

1. *Some preliminary remarks*

Article 8 forms the gateway to Section 2 of Chapter III, i.e. the Section setting out the requirements for high-risk AI systems, the substantive part of regulating these systems. The section continues with seven articles setting out requirements (or, in Article 15, a bundle of requirements) to be met by high-risk AI systems. Article 8 contains two clarifications that serve, especially the first one, to measure the scope of these requirements, an essential task because they are generic requirements that require precision.

This article has changed a lot during the drafting process. The initial proposal only stated:

1. *High-risk AI systems shall comply with the requirements established in this Chapter.*
2. *The intended purpose of the high risk AI system and the risk management system referred to in Article 9 shall be taken into account when ensuring compliance with those requirements.*

* This paper has been written within the framework of the Research Project "Algorithmic tools for citizens and public administrations", funded by the Spanish Ministry of Science and Innovation (PID2021-126881OB-I00).

The entire current paragraph 2 is new, it did not appear in the Commission proposal. In addition, the current paragraph 1 (which was paragraph 2 in the initial proposal) has been enriched with a centrally placed reference to "the generally recognised state of the art in AI and AI-related technologies". These innovations stem from Parliament's amendments, although, as we shall see, they were significantly modified in the final negotiation.

Article 8 corresponds to paragraph 64 of the preamble, which does not provide important details (except for paragraph 2 of the Article).

2. *Criteria to be taken into account in the assessment of compliance with the requirements (paragraph 1)*

It is unavoidable to emphasise, when setting out the requirements of high-risk AI systems, their status as mandates of generic, vague and indeterminate content. They are principles in the sense that they set an objective or trend, the absolute fulfilment of which is almost impossible to achieve and which can be fulfilled in varying degrees or forms. In fact, it may even happen that several of these principles, taken to the extreme, collide with each other so that it is necessary to arbitrate a compromise.

It is not difficult to give examples to prove this assertion. Thus, the risk management system of Article 9 does not seek to eliminate risk (which would be achieved by the "zero option" of not using high-risk AI systems) but rather to reduce it and achieve an "appropriate balance", admitting the existence of a "residual risk" which, reasonably, "is judged to be acceptable".

The quality criteria relating to "Data and data governance" (Article 10) must be "appropriate", and data must be "*to the best extent* possible free of errors". Absolute objectives are not required but relative, "proportionate" ones, so benchmarks are necessary.

Similarly, the record-keeping requirement of Article 12 requires that logging capabilities be in place to achieve an "appropriate" level of traceability.

The transparency required by Article 13 must also be "sufficient" and "appropriate".

Human oversight (Article 14) requires the application of an "appropriate" interface and measures that are "commensurate" with the specific risks and characteristics of the system. It must also ensure that the deployer can take control in a way that is "appropriate and proportionate".

Even more evident is Article 15, which establishes as requirements some concepts that are a maximum objective ("accuracy, robustness and cybersecurity") whose 100% achievement is impossible unless the use of this technique is renounced, so what is relevant is to know what the specific level of requirement is, something that is done employing indeterminate concepts. Thus, an "appropriate" level is required for these concepts.

It is not only that the regulation cannot contain all the technical details, which are necessarily left to other development instruments (regulatory or otherwise), but also that the criteria that determine the degree of exigency in fulfilling the requirements established in Articles 9 and following are undetermined. These criteria are significant decisions by means of which conflicting interests and principles are weighed up against each other. Thus, a stricter compliance requirement (i.e. setting a very low threshold for acceptable risk) means reducing the possibilities for the use of AI or forcing AI to be safer and, at the same time, more expensive. It is the same as in any other field: requiring stringent safety measures in cars makes them safer but also more expensive, which influences the market (by driving out certain operators or products) and consumers (because it can make these items more expensive, even if it makes them safer). The same happens in the housing sector through instruments such as technical building codes, which, in an apparently neutral way, actually imply decisions with social and political relevance by making products and access to housing more or less expensive (depending on the options chosen).

The initial version of the article already spoke of the "intended use" of each system as an element that should be considered "when ensuring compliance with those requirements". It is clear that the intended use of an AI system is very relevant because it determines its possible negative effects. A system that a company uses to decide where to invest is not the same as one that is used to take it into account when deciding who is admitted to a Master's programme. What happens is that this criterion (the intended use) is already used to determine which AI systems are high-risk. It is the fundamental criterion taken into account for Annex III since it is based on the purpose of the AI systems, not on any other characteristic.

Therefore, there is not much that the intended use can contribute to Article 8, considering that it has already been implemented previously.

The Parliament also proposed to incorporate "the reasonably foreseeable misuses" of the AI system as a criterion to be considered, alongside intended use. It is true that when assessing the risks generated by an IA system and the measures to be taken to mitigate them, it should not be assumed that the

system will always be used in the intended direction. It should be borne in mind that the first risk is that it will be used for a different (and negative) purpose, but that this is also possible given the configuration of the system. Similarly, when assessing the desirability of implementing a public policy or enacting a regulatory reform, one should not only take into account the expected or desired effects but also the negative effects that may also occur with some degree of probability, even if they are not desired by the drivers. After all, systems, once deployed, are in the hands of their implementers, who will rationally determine their own behaviour according to their interests, not necessarily according to the preferences of those who have designed the system (who have other incentives altogether). Indeed, other articles in this section already mandate that reasonably foreseeable misuses be taken into account when assessing risks and measures to be taken. This could also have been included in Article 8.1, but it is not essential.

As for the risk management system in Article 9, which is also mentioned in Article 8 (and, moreover, since the original proposal), it is not very helpful. The mention can be considered almost redundant. The risk management system means that the supplier has to justify compliance with the requirements, bearing in mind that there are probably several ways of achieving compliance and that assessing compliance may be difficult or debatable. The risk management system serves, for example, to isolate or identify the risks generated by an IA system so that it is also possible to determine the remedies that need to be put in place to address them and to assess their adequacy.

The risk management system does not indicate what the threshold level is, what the requirements are or how they are to be interpreted but rather, it justifies compliance with them. Its function is, therefore, not the same as that of other concepts mentioned in Article 9, such as the state of the art.

The most relevant novelty in the final wording of Article 8(1) is the mention of the "generally acknowledged state of the art *on AI and AI related technologies*" (these last words were added in the final pact from the text proposed by the European Parliament). The state of the art is a standard generally used in the regulation of different sectors related to technological progress, such as, for example, environmental protection. The state of the art is characterised by the fact that it is variable and evolutionary and does not, in principle, admit regression (the economy or available resources may regress, but not knowledge). The subjects, including economic operators in particular, are obliged to be aware of the state of the art and assume the cost of keeping up with it, and applying it to their products. The state of the art operates similarly to the "progress clause" or "innovation clause" in concession

contracts. The concessionaire must assume the cost of complying, in its services, not only with the conditions and requirements that were customary at the beginning of the concession term, but also with those that have become customary subsequently as a result of technological innovation. This is a cost assumed by the concessionaire as part of the risks transferred to it.

Similarly, in environmental matters, companies must use the most appropriate technologies in accordance with the state of the art, and only pollution that cannot be eliminated by using these technologies will be acceptable. Solutions based on obsolete and, therefore, more polluting technologies are not acceptable.

It is clear that, when assessing whether an AI system meets the requirements set out in this section, the performance and possibilities offered by the available techniques, i.e. those that are part of the state of the art, must be taken into account. Risks resulting from the use of obsolete technology, which is not in line with the state of the art, cannot be accepted.

Determining such a state of the art certainly requires technical support, the application of expertise that not only AI system providers but also notified bodies and supervisory authorities need to have at their disposal.

Regardless of the knowledge required to appreciate the state of the art, some conceptual issues need to be clarified. State of the art is a common and accessible standard of knowledge. It does not equate to the cutting edge of knowledge, to still experimental solutions or at the frontier. With that single qualification, the state of the art includes all known technologies or developments.

The state of the art must always be considered, but is it necessary to always use the most advanced technologies? I believe that the reference in Article 8(1) to the state of the art should be combined with the other concepts used in this section when regulating the different requirements, and which I have quoted *above*, on "appropriate", "proportionate", etc. It is not always necessary to apply the same standard. The standard must, in each case, be proportionate to the circumstances. However, in assessing that proportionality, the starting point or benchmark should be the technological level offered by the state of the art, not a lower one.

A critical aspect of the state of the art concept is its evolutionary character. The state of the art refers to the evolution of science and, as a consequence, of technology. This is why it is used in cases with a long-lasting legal relationship (such as concessions) to ensure that the standards adapted to the state of the art are met in each case, not those that were appropriate at the outset.

This also means that the conformity assessment of AI systems must be adapted to the circumstances. It is true that AI systems will not usually be in use for a very long time, nor can we assume that there will be major changes in AI technology, but the system will have to meet the requirements of this Section according to the state of the art at the time. If significant advances are made, they must be incorporated, or the system must cease to be used. A system may initially meet the requirements but subsequently cease to do so.

The text added by Parliament did not only include the reference to the state of the art, which is finally retained in the adopted text but also referred to the various implementing instruments provided for in the Regulation, such as harmonised standards or guidelines. The deletion of this reference does not seem significant to me because these instruments retain their binding effects and must be considered to define the specific scope of the requirements because that (and no other) is their function.

Moreover, the reference also goes in the opposite direction because the state of the art will undoubtedly have to be taken into account in the drafting of guidelines and harmonised standards, whose function is essentially to translate it into concrete requirements.

3. *Compliance with the requirements regarding to products to which requirements of the Union harmonisation legislation listed in Section A of Annex I apply (paragraph 2)*

This section deals with how compliance with the requirements should be verified for high-risk IA systems that are high-risk because of their relationship with products already subject to European risk prevention regulations. The aim is to coordinate both verification procedures in order to avoid bureaucratic overload.

This paragraph, which did not appear at all in the Commission's proposal, has undergone many changes, from the wording proposed by the European Parliament to the text adopted in the final agreement.

One of the fears of the Regulation is that it will lead to excessive bureaucracy and that the red tape involved in bringing an AI system or an AI product into service will paralyse innovation. In view of the criticisms received, including that of the Draghi report, it does not appear that this concern has led to sufficient remedies, at least in the view of the critics. In any case, the Regulation attempts to avoid duplication of control procedures in several precepts, including this one, by making it easier for existing procedures also to serve to verify compliance with the requirements established in this Regulation for AI systems.

In summary, paragraph 2 states that AI systems attached to products that are already subject to EU regulation must ensure that they comply with all the requirements established by the different standards (in this case, the one that regulates the product and the one that regulates the AI system associated with it). Furthermore, it tells us that when the regulation to which this product was already subject before being "enriched" with an AI system is a regulation of the "new legislative framework" (standards listed in Annex II, Section A), the supplier may choose and accumulate in a single procedure (the one regulated in that standard included in the new legislative framework) the submission of documents and the verification of compliance with all the requirements.

The article recalls (and the preamble reinforces this idea) that it is normal for a product or service to be subject, by virtue of its purpose, to several European regulations and that all of them must be complied with in order to be put into service. This applies in particular to high-risk systems in Article 6.1, i.e. those which are high-risk because they are connected to products already subject to one of the regulations listed in Annex II because of the associated risks . These systems are already subject - by definition - to a European regulation that aims to prevent the risks they already produce outside of IA.

The amendment proposed by the Parliament was mandatory (leaving no choice to the provider) and avoided duplication by directly merging the two control procedures. This wording referred only to the standards in Annex II, Section A (as in the text finally adopted). If the product was subject to one of these standards, and this *includes the AI component* (i.e. it has been modified or updated to also refer to AI risks), the European Parliament's proposal stated that the requirements established by the Regulation for high-risk AI systems would be deemed to be fulfilled by the application of this other European regulation which already includes the AI component. Duplications are avoided by direct imposition of the regulation itself. Moreover, in the case of the standards in Annex III, Section A, which do not (yet) contemplate the requirements relating to the IA applied to the product or service in question, what this proposal mandated is that the standard be adapted to include them, in which case there would also be this simultaneous review of the two types of requirements: those inherent to the regulated product and those relating to the IA system added to it (and which is high risk). In practice, prior to the adoption of the Regulation, the authorities in charge of applying these European standards for risk products were already reviewing also, where applicable, the risks arising from the AI systems attached to or incorporated in them.

The text finally adopted no longer aims to ensure that the regulatory standards for products include, alongside the risks inherent to them, the risks arising from the inclusion in them of "embedded" high-risk systems. On the contrary, it is considered normal for a product to be subject to several regulations simultaneously, and the supplier must ensure, before putting it into circulation, that it complies with all of them. As stated in paragraph 64, citing the Commission Communication entitled 'Blue Book' on the implementation of European product legislation of 2022, it is not unusual, but rather the general rule, for a product to be subject to 'more than one piece of EU harmonisation legislation'. The different pieces of legislation deal with different things and must be subject to "simultaneous and complementary application".

This finally approved text, which no longer aims to ensure that the product incorporating an IA system is subject to a single standard, simply *allows* the supplier to cumulate the verification of compliance, in whole or in part, in a single procedure, the one generally applicable to that product and regulated in one of the standards of Annex III, Section A. Flexibility is simply required: "flexibility with regard to operational decisions on how to ensure compliance of a product that contains one or more AI systems with all the applicable requirements of that Union harmonised legislation in an optimal manner. That flexibility could mean, for example, a decision by the provider to integrate a part of the necessary testing and reporting processes, information and documentation required under this Regulation into already existing documentation and procedures required under existing Union harmonisation legislation based on the New Legislative Framework and listed in an annex to this Regulation. This should not, in any way, undermine the obligation of the provider to comply with all the applicable requirements".

This means that the supervisory authority in charge of implementing this other regulation will at least have to check compliance with the IA Regulation. Similarly, if the intervention of a notified body is necessary, one and the same body should also be competent to assess compliance with the IA requirements. If the intervention of two authorities and two bodies were necessary, little progress would have been made in eliminating duplication.

None of this, of course, takes away the powers of the market surveillance authority, which will be able to take all the measures required by this Regulation regardless of which authority intervened before the system was placed on the market.