Chambers Ireland
100
YEARS OF ADVANCING
BUSINESS TOGETHER

13 CLIMATE ACTION

5 GENDER EQUALITY

11 SUSTAINABLE CITIES AND COMMUNITIES

9 INDUSTRY, INNOVATION AND INFRASTRUCTURE

8 DECENT WORK AND ECONOMIC GROWTH

Public consultation on the proposal for a Directive on adapting non-contractual civil liability rules to Artificial Intelligence

Submission by Chambers Ireland

January 2023

About Chambers Ireland

Chambers Ireland is an all-island business organisation with a unique geographical reach. Our members are the Chambers of Commerce in the cities and towns throughout the country – active in every constituency. Each of our member Chambers is central to their local business community and all seek to promote thriving local economies that can support sustainable cities and communities.

Our Network has pledged to advocate for and support the advancement of the United Nations Sustainable Development Goals (SDGs). Accordingly, we use the Goals as a framework to identify policy priorities and communicate our recommendations. We have a particular focus on five of the goals encompassing decent work and economic growth (SDG 8), sustainable cities and communities (SDG 11), Gender Equality (SDG 5), Industry, Innovation and Infrastructure (SDG 9) and climate action (SDG 13).[1]

In the context of the current consultation, Industry, Innovation and Infrastructure (SDG 9) is the most relevant goal. Chambers Ireland has contributed to the national conversation on artificial intelligence, having previously submitted our views regarding the development of a National Artificial Intelligence Strategy in 2020[2] and worked with our European Colleagues on our Submission to the European Commission[3]. AI tools will increasingly be used to automate tasks that people already do, aiding productivity and helping them focus on the important elements of their work. When people are choosing to apply these tools to their work it is essential that they are responsible for their choices. Our submission is written with the overriding point that any regulatory framework regarding AI does not absolve them of their responsibilities to apply these tools wisely and appropriately. As such, the context, use and environment for employing AI systems is critical when assigning liability where fault occurs.

[1] The Chambers Ireland SDGs. Available at: https://www.chambers.ie/policy/sustainable-development-goals/chambers-ireland-sdgs/
[2] ai-strategy-submission-3-chambers-ireland.pdf (enterprise.gov.ie)
[3] https://www.eurochambres.eu/publication/eurochambres-contribution-to-the-commission-consultation-on-adapting-liability-rules-to-the-digital-age-and-artificial-intelligence/

Contents

Key Points

- We welcome the introduction of a non-interventionist approach to AI regulation.
- A fault-based liability regime is preferable to a strict liability regime for assessing liability regarding AI systems.
- When assigning liability, emphasis should be on use and environment, not just the AI system itself.
- Any disclosure of information must consider the type of information involved; i.e the information it is trained on, the metadata of the AI system, how it operates, or how specific, effective or ineffective it is.
- Any regulation of AI ought to ensure innovation is not stifled, while at the same time facilitating access to remedies for injured parties.
- Clarification is required regarding the disclosure of information, establishing a causal link, and determining what constitutes a ‘high-risk’ AI system.

The perspective of Chambers Ireland regarding the regulation of Artificial Intelligence

There are two bad regulatory environments for AI. Firstly, regulation could be used to protect incumbent industries which are at risk of automation from the forces of competition particularly in non-traded sectors. Support for such a regime could emerge where sectors and industries which have previously been protected from external competition suddenly have to accommodate new entrants to the market, whether they be home-grown competitors or competitors which originate from abroad. The temptation may be to erect barriers to competition which will be defended using fears about AI technologies. Ultimately this will result in our domestic economy's productivity diminishing relative to competitor nations.

Secondly, there is the risk of making categorical errors regarding the nature of the AI technologies which are available to us and creating a regulatory regime that is based on a mischaracterisation of what these tools are, leading to potential local innovation in the field of AI being further incentivised to offshore.

Any regulation of AI needs to ensure that it does not stifle innovation and leave the innovational gains to be achieved in other jurisdictions where overregulation does not occur. It should not be unduly interventionist or constitute a regulatory overreach against the defendant. At the same time, it should not only ensure that remedies are available, but that they are accessible to those requiring them.

The increased use of AI in the future is inevitable. In our submission from 2020, we emphasised that exaggerated claims regarding AI are a hindrance when it comes to the appropriate regulation of the sector. A realistic appraisal of AI is necessary when designing policies affecting the use of AI. The real gains from AI will be in its utilty in providing non-critical services, rather than high-risk services. We anticipate that in future its use in automating non-critical services will increase.

One of our main concerns regarding AI is the fact that its application and the decisions made by the person using it – not the AI tool itself - in different environments is of primary relevance. Simply because AI is involved does not automatically mean that the AI tool is at fault. Nor does it mean that the provider of the tool is solely at fault. Determining the appropriate use and environment for the AI tool is especially important when apportioning liability and figuring out who exactly is responsible for the injurious action.

There are two significant risks that emerge from regulating AI tools as a special class for liability. Firstly, there is the risk of over-regulation that stymies the development of AI tools within the single market. This is likely to have two second order effects; there will be fewer high tech jobs within the single market as other jurisdictions will come to dominate this economic space, and that when single market persons use AI tools that are developed using data that comes from non-single market jurisdictions, then these tools are less likely to be as effective when used by single market consumers than they will be when used by individuals from their home markets. Alternatively, there is the risk that by making AI a special class of tool that has its own unique liability rules there is the risk that such rules could become unenforceable, either because of the complexity of the technical arguments that are needed to determine where liability lies, or because legislation does not keep pace with the state-of-the-art practice in the field. Should that happen, then there is risk that a de facto zero-liability regulatory regime could emerge. The reification of AI could alter behaviours among developers and users that leads them to use AI tools not because they are the best tools for that particular task, but because the practical barriers to making claims against the users/developers are such that the use of an AI tool will effectively indemnify them against future liabilities.

Our call for effective regulation of AI in 2020 aligns with our support for the introduction of an AI Liability Directive at EU level. We welcome the attempt to harmonise liability rules for AI. Currently, no option exists to seek compensation for damage by AI systems under the Product Liability Directive, so the proposal is necessary. Creating certainty in the regulation of AI will help both those developing the technology and those using it, as the absence of a tailored legislative framework has been a real concern for both cohorts. We urge all legislators, at the national and European level, to treat AI related liability as other business or personal tools are treated. Ultimately humans make decisions; sometimes humans may choose to automate a particular process, but if they use the wrong tool to automate that process, or if they automate a process that should not be automated, then the liability should lie with the person/institution that made that decision. Alternatively, if their decision was based on inaccurate information about how or when to use a tool then the developer should be responsible for supplying an inappropriate tool to the consumer.

Should AI tools be treated in a way that is different to the use of other tools that persons use, then there is a significant risk that the technology obscures the decision-making process that led to a particular tool being used inappropriately in a particular circumstance. This may mean that individuals that would otherwise be entitled to court remedies may find it difficult to have their rights recognised.

One of the aims of the Commission's 2020 paper was to ensure that any new regulatory framework for AI should be effective in achieving its objectives while not being excessively prescriptive so that it could create a disproportionate burden, especially for SMEs.[4] Though not without criticism, we welcome the proposal taken by the Commission to use a risk-based approach to avoid unnecessary regulatory burdens and ensure it is proportionate.

Notwithstanding the criticisms outlined in our submission regarding the proposal, this is a step in the right direction. We anticipate that, subject to further changes, this will bring legal certainty to SMEs and businesses, in laying down concrete liability rules outlining redress for parties injured by AI technologies.

[4] https://commission.europa.eu/system/files/2020-02/commission-white-paper-artificial-intelligence-feb2020_en.pdf

The perspective of Chambers Ireland regarding specific aspects of the proposed AI Liability Directive

One of our chief concerns regarding AI systems is an over-interventionalist approach to regulation. Regarding this concern, we welcome the fact the Commission has proposed a fault-based liability system, rather than a strict liability system. In our view, a strict liability system runs the risk of incorrectly favouring claimants at the expense of defendants. Nonetheless, fault-based systems require sufficiently clear criteria to establish and assess liability where damage has occurred. Elements of the current proposal require clarification if it is to be effective.

One of our chief concerns of the proposed Directive is the lack of clarity regarding (a) the burden of proof, (b) establishing causal links, and (c) the disclosure of evidence. There are other formulations in the proposal which are not adequately defined. Much of our submission focuses on bringing extra clarity to the proposed articles. For example, it is unclear how the "relevant" requirement for evidence will be relied on when assessing a claim for damages. Similarly, it is unclear to what extent evidence is "necessary" and "proportionate" when assessing a claim for damages. These are open-ended requirements that leave interpretation up to the discretion of a judge in a court of law. This will not aid uniformity and if left inadequately defined, risks regulatory fragmentation and confusing claimants and developers alike.

Regarding specific articles:

Disclosure of information

Under the proposed Directive, courts can order AI providers to disclose information about their systems to assess liability. However, there is a relatively-simple get-out-of-jail clause; defendants may prove compliance with risk management measures in the AI Act to remove the requirement to disclose information. As stipulated elsewhere in our submission, an appropriate level of risk needs to be taken into account when assessing liability; this is because AI tools are probabilistic by their nature and consequently their failure rate needs to be considered. We have the concern that a self-certified compliance regime will not be sufficient to ensure that all risks related to the use of a particular AI tool are taken into account. Article 9 of the AI Act requires that risk assessments are carried out for all high-risk AI tools, and a concern that Chambers Ireland has relates to the “reasonably foreseeable” qualification on risk assessment. Should the risk assessment element devolve into a merely procedural process, or becomes equivalent to a voluntary guideline, which only gets updated when a known issue has become established, it is not likely to have the intended effect. There is a significant risk that allowing the defendant to determine what is reasonably foreseeable will damage the effectiveness of the Directive, as in effect it leaves liability to be determined by the defendant. We have seen in other sectors that the auditing of best practice is often inadequate and it is reasonably foreseeable, in this context, that the auditing of AI tool risk assessment could become a paperwork exercise. Should the presence of a risk assessment process be sufficient to exclude a defendant from liability, and also the obligation to reveal any system information, then it is likely to obscure circumstances where a risk assessment was inadequate.

It is also not clear what type of information is being requested in relation to the AI system itself. In this instance, three types of information must be distinguished: a) the information the system is trained on; b) the metadata regarding the expected operation of the system; and c) the actual data that was inputted to the system at the time of operation. Information as to how specific the system is will be easy to ascertain. However, understanding how it operates/does its job is more difficult. Accordingly, we are concerned with the phrasing of Article 3(1) that information or evidence deemed as being related to the “specific high-risk AI system that is suspected of having caused”, because this may mean that the damage is not viewed in the appropriate context. In line with our point about the crux of an issue being the *environment* AI systems are adopted in, and their *use* in those environments, we believe this does not consider that AI systems cannot be viewed independent of the circumstances of their use when investigating damage caused or ascertaining liability connected to that damage.

To aid certainty for providers, the proposal should lay down timeframes and parameters for complying with a disclosure request, as currently, this is not mentioned in the proposal. For example, there could be a lot of private information in the data that is used to train the AI-tool. The developers may not be able to share all that data because releasing that data may reveal the private information of individuals. Pseudo-anonymisation is so hard that it just doesn't work, and gaining access to such datasets may lead to that information being linked to particular people. Alternatively, there could be so much information that not all of the information that was used in the training of the tools is available, or even stored, and so it may be impossible to return a complete set of training data to the applicant.

Finally, ‘plausibility’ as a test for requesting disclosure is problematic. It remains to be seen how national courts will interpret the plausibility of a claim, so this may invoke regulatory fragmentation. Further clarification is required to specify what constitutes a plausible claim.

<u>‘High-risk’ AI systems</u>

‘High risk’ AI systems are the main focus of the proposed Directive. Unfortunately, determining what is high risk is difficult, and consideration should therefore be given to a set of criteria for determining what meets the ‘high risk’ definition. In conjunction with meeting the definition, the AI system needs to be related to eight areas identified in the proposed Directive. i.e among others, identifying natural persons, managing critical infrastructure, education + training, employment, private + public services, law enforcement, migration, and administration of justice. We welcome this list, however the absence of set criteria for determining what constitutes a ‘high-risk’ AI system leaves room for confusion for claimants and defendants.

<u>The burden of proof</u>

It is important to strike the correct balance between ensuring access to a remedy, while ensuring the Directive does not stifle innovation by being overly-interventionalist or placing a disproportionate burden on developers. The high burden of proof runs the risk of rendering the

Directive largely ineffectual for claimants. This potentially excludes those who may have been injured and should be entitled to a remedy.

The burden of proof as set out in the proposal is too high and risks striking regulatory imbalance. The wording of the Directive means that there is a requirement for the claimant to: (1) demonstrate non-compliance with (2) a duty of care intended (3) to protect against the risk of damage. In our view, this is a very high burden of proof for claimants.

Presumption of a causal link in case of fault

The conditions for the presumption of a causal link pursuant to Article 4 lack clarity and there is an overreliance on the term 'relevant'. Article 4(1) also does not adequately set out how the defendant (i.e. the provider) can be identified. It should be considered that AI systems make probabilistic decisions based off previous data. Therefore, they are not certain and error rates have to be accounted for when deciding to use them. How frequently the AI system will be wrong, and what degree of error is tolerable will be closely tied to its use case. The risk is that the developer will be wrongly identified as the defendant, rather than the person who employed the tool in the wrong environment. This should be considered when setting the criteria for establishing a causal link, and precise wording is required to determine this. In addition, more consideration should be given to the extent to which providers are liable when it comes to advising the how/when/where/who of using an AI tool.

Article 4(3) needs to account for the fact that it is the insertion of data which is used in the inappropriate circumstance or use case, that is likely to have contributed to the contested output that should be considered a violation of the duty of care by a user, rather than the introduction of non-relevant data.

The presumption of a causal link is not available where the causal link is established based on information obtained through disclosure. This has the potential to create problems down the line, making it difficult for claimants to avail of a remedy.

Article 4 also states that the defendant/party held liable should be the provider of an AI system that has not previously complied with one of the "duties of care" listed in the AI Act proposal. This is a problematic formulation because in many cases it is difficult to determine who the provider is. As an example, there may be cases where the injured party is in contact with the final user of the AI system, instead of the provider. A mechanism that allows the defendant to be determined at an early stage should therefore be considered.