

Expert explainer

Christiane Wendehorst

Professor of Civil Law,
University of Vienna

AI liability in Europe: anticipating the EU AI Liability Directive



September 2022

Introduction

This explainer is intended to provide those interested in the EU AI Liability Directive and the AI Act with legal context and analysis of how and where product and fault liability could be leveraged. It provides five reasons for EU legislators to act, with three illustrative scenarios and three policy options to address under-compensation, and a commentary on AI liability beyond traditional accident scenarios.

It will be particularly useful for EU, UK and global policymakers who are interested in the progress of the AI Act, and in understanding how liability could support a more effective legal framework for AI.

The European institutions have been addressing the challenges of the digital economy at a breathtaking pace. This year has already seen several major pieces of legislation being adopted, including the Digital Services Act, the Digital Markets Act and the Data Governance Act. A number of others have reached advanced stages within the legislative procedure, including the draft AI Act, draft Machinery Regulation and draft Data Act, to name but a few.

With all this legislative activity, the EU seems currently to be the most dynamic region in the world when it comes to regulating the digital sphere. But there is one legislative project that has repeatedly been postponed and now seems to bring up the rear: new liability legislation for emerging technologies, in particular artificial intelligence.

Two legislative proposals have now been announced for 28 September 2022: a revised Directive on product liability, replacing the existing Product Liability Directive 85/374/EEC, and a brand new AI Liability Directive, without precedent in EU legislation, which is the focus of this explainer.

What does liability law bring to AI regulation?

Liability law has the potential to provide answers to important questions regarding the legal consequences of harms caused by AI systems. Who is liable when an AI system fails or malfunctions, and the role of human decision-making is opaque? There are many actors involved in the chain of events leading to a potential instance of harm: designers, manufacturers, data providers, deployers, employees working with the AI and so on.

The removal of human decision-making from systems that risk harm raises important legal challenges for existing frameworks, as well as ethical and societal challenges. It is in the interest of citizens, businesses and regulators that we get liability for AI right. We cannot make AI work for people and society without it.

Timeline

Preparations for a new piece of legislation on AI liability have been underway for some time. They began with the 2017 legislative resolution of the European Parliament on Civil Law Rules on Robotics,¹ which became famous for recommending the attribution of 'electronic personhood' to the most advanced robots. This was met with heavy criticism throughout Europe and was subsequently dropped.

In 2018, the European Commission established an Expert Group on Liability and New Technologies, comprised of two separate Formations. One of these, the New Technologies Formation, submitted an Expert Group Report on liability for AI² and other emerging digital technologies in 2019 (co-authored by the author of this explainer), which served as a basis for further EU activities. In 2020, the Commission's Report³ on the safety and liability implications of AI, the Internet of Things (IoT) and robotics still left more or less all approaches to liability open.

1 European Parliament. (2017). *European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL))*. Available at: https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html

2 Expert Group on Liability for New Technologies. (2019). *Liability for Artificial Intelligence*. European Commission. Available at: https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=63199

3 European Commission. (2020). *Report on the safety and liability implications of Artificial Intelligence, the Internet of Things and robotics*, COM/2020/64 final. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020DC0064>

Meanwhile, the European Parliament had become active again, putting further pressure on the Commission. In October 2020, the Parliament passed a fully-fledged proposal for a Regulation on AI liability. It took the Commission another year to respond with the opening of a public consultation,⁴ which ran from October 2021 until early January 2022. The results of that consultation have now been published,⁵ but the exact details of what will be proposed in autumn 2022 remain uncertain and unconfirmed.

What seems increasingly certain, however, is that the Commission is now heading towards a Directive rather than towards a Regulation. The basis for this decision is the high degree of overlap and interaction between national regimes of tort law, which differ greatly. To derogate all these regimes by way of a Regulation, but restricted to harm caused by AI, could lead to an unacceptable level of friction and inconsistency across the EU.

Types of liability

Fault liability: liability based on the defendant's fault (i.e. intent or negligence) in causing the harm; negligence could consist, e.g., in failing to apply due diligence in designing, deploying or monitoring an AI system.

Product liability: liability of the producer (or, under certain conditions, other players such as importers) for harm caused by defective products; this type of liability does not require the defendant's fault and was harmonised throughout the EU by Directive 85/374/EEC.

Strict liability: liability not requiring the defendant's fault; this still covers a spectrum, up to liability for mere causation of harm without any further requirements (but usually with a defence of *force majeure*).

Vicarious liability: liability for the conduct of others, such as auxiliaries.

4 European Commission. (2021). *Civil liability – adapting liability rules to the digital age and artificial intelligence*. Available at: https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12979-Civil-liability-adapting-liability-rules-to-the-digital-age-and-artificial-intelligence_en

5 European Commission. (2021).

Types of EU legislation

Directive: EU legislation that is binding, as to the result to be achieved, on each Member State to which it is addressed, but leaves to national authorities the choice of form and methods; a Directive needs transposition by the national legislatures.

Regulation: EU legislation that has general application, i.e. is binding in its entirety and directly applicable in all Member States without transposition by national legislatures.

Why do we need a new regime of AI liability?

AI technologies present novel challenges to existing frameworks, the current legal frameworks are fragmented and incomplete, and there is a dependency on the national liability regimes of Member States which undermines the objective of creating a level playing field for businesses across the internal market. It may therefore come as a surprise that the Commission took so long to take concrete action towards establishing new rules for AI liability.

However, regulating AI liability is more complicated than regulating the issues addressed by the AI Act and other forms of digital regulation, precisely because Member States already have very sophisticated and longstanding liability rules, so it is not obvious from the outset that action by the EU institutions is required.

However, there are five clear reasons for the European legislator to act on AI liability, which are:

1. Avoiding under-compensation for injured parties

The main argument for adapting liability rules to AI, or introducing new liability regimes, is that it would help prevent under-compensation of injured parties where the harms were inflicted by AI systems. Under-compensation may result from the absence of an appropriate legal response and/or from the legal process of seeking compensation for AI-related harms becoming unduly difficult or expensive.

This should not be taken to imply that there is a unanimous view on the 'right' level of compensation. Scholars and policymakers have been struggling for centuries with the question of to whom losses should be attributed (e.g. whether the primary criterion is wrongful conduct, benefit derived, degree of control or being the cheapest cost avoider). However, most would agree that the adoption and use of AI systems in society should, at the least, not leave injured parties worse off than before with regard to compensation.

2. Enhancing enforcement of the AI Act and similar legislation

When the AI Act proposal was presented in April 2021, it came as a surprise to some that, unlike the General Data Protection Regulation (GDPR), the AI Act takes a traditional product safety approach to regulation. This means that the proposal includes a list of essential requirements which certain high-risk AI systems have to meet in order to be placed on the market, ranging from data governance to human oversight, transparency and robustness. It also includes a number of obligations for AI providers and others in the supply chain, including the deployer (normally a business or public authority).⁶

What it does not include, however, is individual rights on the part of those affected by the use of AI (e.g. citizens, consumers), such as the right to claim damages where harm has been caused. There is therefore no private enforcement in the AI Act proposal itself. This is something that could potentially, at least to some degree, be fulfilled by the AI Liability Directive. The Ada Lovelace Institute, in its policy briefing,⁷ recommended including ‘affected persons’ within the AI Act, defined as those natural or legal persons who are ultimately affected by the deployment of an AI system.⁸

3. Increasing public trust in new technologies

Leaving aside any actual under-compensation of injured parties, the introduction of a new regime of EU-wide AI liability has important symbolic value, within the EU and globally. It demonstrates that the EU is acting to protect EU citizens against new and potentially harmful technologies. This action is likely to increase public trust in new technologies, in particular AI. This is all the more important as the public perception of new technologies is often fuelled by science fiction

6 For an expert legal opinion on the EU AI Act’s product safety approach, see: Edwards, L. (2022). *Regulating AI in Europe: four problems and four solutions*. Ada Lovelace Institute. Available at: <https://www.adalovelaceinstitute.org/report/regulating-ai-in-europe/>

7 See: Ada Lovelace Institute. (2022). *People, risk and the unique requirements of AI*. Available at: <https://www.adalovelaceinstitute.org/policy-briefing/eu-ai-act/>

8 A similar recommendation was made by the author of this explainer. See: Wendehorst, C. (2021). *The Proposal for an Artificial Intelligence Act COM(2021) 206 from a Consumer Policy Perspective*. Austrian Federal Ministry of Social Affairs, Health, Care and Consumer Protection. Available at: https://www.sozialministerium.at/dam/jcr:750b1a99-c5af-47bd-906a-7aa2485dabbd/The%20Proposal%20for%20an%20Artificial%20Intelligence%20Act%20COM2021%20206%20from%20a%20Consumer%20Policy%20Perspective_dec2021__pdfUA_web.pdf

rather than facts, leading to widespread fears of AI systems killing both jobs and people.

EU legislation needs to make sure that the AI systems brought to market are safe and trustworthy, but the public also expects another important 'safety net' in the form of liability for cases where harm nevertheless occurs. The European Parliament's proposal has created additional pressure on the Commission, potentially increasing public expectation of EU-wide action.

4. Ensuring a level playing field and innovation-friendly climate for businesses

Somewhat paradoxically, some of the calls for a new regime of AI liability are primarily driven not by concerns about injured parties and their right to compensation, but rather by concerns about innovation and the regulatory environment for businesses. Many of those who want to see a pro-innovation regime across the EU worry that national legislatures and/or courts will act in ways that create unnecessarily strict and extremely divergent rules for AI liability across the internal market.

This could stifle innovation and would seriously hamper efforts towards the ambition of a Digital Single Market. If the future AI Liability Directive were to follow a maximum harmonisation approach (which remains to be seen), i.e. if it were to prohibit Member States to maintain or introduce stricter liability regimes themselves, this would be able to protect businesses from excessive liability and legal uncertainty under national laws.

5. The 'Brussels effect'

Last but not least, there is what has come to be known as the 'Brussels effect', a term first coined in 2012 by Anu Bradford, mirroring the perceived 'California effect' within the USA.⁹ The term expresses the EU's role as a global leader when it comes to creating regulatory

⁹ Bradford, A. (2012). *The Brussels Effect*. Columbia Law School. Available at: https://scholarship.law.columbia.edu/faculty_scholarship/1966/

concepts and principles, and the ways in which these go on to shape the development of laws outside of the EU. In a broader sense, the term can be understood as referring to the competitive advantage of being the regulatory 'first mover', significantly increasing the chances that other regions of the world will be directly inspired by EU rules.

Successfully conceptualising and regulating the digital sphere therefore has broader strategic importance for the EU, and this includes developing a regime for AI liability. Of course, we should also not forget that 'Brussels' does not simply act as one unit, as it consists of a multitude of different entities and individuals, which are competing for visibility, influence and success.

Understanding under- compensation: Why might AI leave injured parties worse off?

According to the 2021 public consultation document,¹⁰ AI systems have a number of specific properties which could make it hard for injured parties to get compensation. The Commission states that ‘the lack of transparency (opacity) and explainability (complexity) as well as the high degree of autonomy of some AI systems’ could make it too difficult for injured parties to prove that a product is defective, or that there is fault at play, and to prove the causal link with the damage.

In order to properly evaluate the Commission’s claim it can be useful to analyse some different scenarios and paradigm examples.

Scenario 1: fully self-driving vehicles

Self-driving (or autonomous) motor vehicles are normally subject to special liability and/or insurance schemes.

In the vast majority of EU jurisdictions, the owner of any motor vehicle is strictly liable for any death, personal injury or property damage thereby caused. There is no requirement of fault on the part of the owner or on the part of the driver. This is coupled with mandatory liability insurance, and according to the Motor Insurance Directive 2009/103/EC (as amended),¹¹ the injured party has a direct claim against the insurer.

10 European Commission. (2021). *Civil liability – adapting liability rules to the digital age and artificial intelligence*. Available at: https://ec.europa.eu/info/law/better-regulation/have-your-say/initiatives/12979-Civil-liability-adapting-liability-rules-to-the-digital-age-and-artificial-intelligence_en

11 European Commission. (2009). *Motor insurance – Directive 2009/103/EC*. Available at: https://ec.europa.eu/info/law/motor-insurance-directive-2009-103-ec_en

The UK is one of the very few countries that still rely on fault liability for road traffic accidents, but the Automated and Electric Vehicles Act 2018¹² ensures that insurers will cover any harm caused by self-driving vehicles, even where the owner is not at fault. In essence, this means that an injured party will always be compensated, regardless of whether the offending motor vehicle was powered by AI or driven by a person.

There is thus no under-compensation with regard to accidents caused by self-driving vehicles, and no injured party is worse off than before. Most of the liability debate for self-driving cars is really about the fair distribution of the associated insurance burden. Should this financial burden continue to rest with the owner, or should it shift to the manufacturer if most accidents with fully self-driving cars are caused by faulty design?

Scenario 2: autonomous AI-enabled devices – lawnmowers or cleaning robots

However, not all autonomous appliances capable of causing harm are subject to special liability or insurance schemes, in the way that self-driving cars are. Consider the case of an autonomous lawnmower that severely injures the feet of child playing in a garden, or an autonomous cleaning robot operating in a public space, which knocks over and injures someone in its path.

If a human gardener drove a lawnmower over a child's feet, or if a human cleaner knocked a person over, it would normally be very straightforward for the injured party to demonstrate negligence. As a result, the gardener or cleaner would be liable for damages, and potentially their employer too, depending on the relevant regime of liability for auxiliaries (vicarious liability).

With an autonomous appliance, things are much more difficult as the robot (unlike the human gardener or cleaner) cannot itself be liable, and as many courts refuse to apply the rules of vicarious liability (which might apply to the gardener's or cleaner's employer) by analogy to cases involving the deployment of machines.

12 UK Government. (2018). *Automated and Electric Vehicles Act 2018*. Available at: <https://www.legislation.gov.uk/ukpga/2018/18/contents/enacted>

The vast majority of legal systems would primarily provide for fault liability in this case, meaning that the owner of the appliance is liable only if it breached a duty of care. The owner may have breached a duty of care by deploying the robot for a task it was not designed for, failing to monitor the robot's operations properly or failing to provide proper maintenance, for example through software updates.

But what if there is no such breach, or the injured party is not able to prove there was a breach or that there was a causal link between the breach and the damage? Another route to compensation under European jurisdictions is through the producer's liability for defective products (product liability). But what if the injured party cannot prove that the robot was defective and that the defect caused the accident?

This is the textbook situation in which AI-related harm may result in under-compensation, when compared with the analogous but non-AI case.

Scenario 3: credit scoring AI

Things get even more difficult in cases of pure economic loss (i.e. financial losses not linked to personal injury, property damage or similar harms) or pure non-economic harms (i.e. non-financial suffering not linked to personal injury). Consider the denial of credit to a banking customer, due to a flawed AI system for credit scoring, which consequently forces the customer to sell their business at below market price in a temporary and acute liquidity crisis. As another example, consider the denial of credit on the basis of data points that strongly correlate with the applicant's ethnicity (amounting to discrimination), whether or not that also results in financial loss.

Most of the relevant cases, such as AI systems for credit scoring, recruitment decisions or calculating insurance premiums, will occur in a contractual or pre-contractual setting, and they will most likely concern non-embedded AI (i.e. standalone software marketed as such). Under many European legal systems, demonstrating liability in such settings requires that the bank itself was at fault (e.g. through the governance of its Board) or that it faces contractual liability for the wrongful conduct of its human employees.

However, as has been explained in the context of scenario 2, courts do not generally apply the rules of vicarious liability by analogy to malfunctioning technical systems. Where they do not, the bank in this scenario will not be liable if it can demonstrate that it bought the credit-scoring AI system from a recognised company and has fulfilled all obligations with regard to proper deployment. In contrast with accidents caused by autonomous devices, product liability is not an option within this scenario as producers are normally liable only for personal injury or property damage.

This means there may be a serious ‘accountability gap’ within liability law, where AI is deployed for tasks traditionally fulfilled by humans and where the injury consists of pure economic loss or pure non-economic harm. The operators of AI systems, who otherwise would have been strictly liable for the wrongful conduct of human employees, could ‘hide behind the AI’ as long as they can demonstrate due diligence, such as with regard to monitoring and maintenance.

Three policy options to avoid under-compensation for accidents

The debate at the EU level is clearly focusing on ‘accident scenarios’, i.e. on the type of scenarios illustrated in the previous chapter by scenario 1, and in particular, scenario 2.

The policy options mentioned by the Commission in its public consultation fall into two broad categories: strict liability on the one hand; and proof and procedural-related options on the other, potentially concerning either product liability, fault liability or both.

Strict liability and/or mandatory insurance

The most far-reaching solution would be a framework of strict liability and/or mandatory liability insurance, i.e. establishing the same liability regime for autonomous AI-enabled devices (see scenario 2) as already exists with regard to fully self-driving vehicles (see scenario 1). The European Parliament’s proposal did suggest strict liability for certain ‘high-risk AI systems’, but without indicating which AI systems could be qualified as high-risk.¹³

Strict liability would mean that the operator of the autonomous device, such as a lawnmower or cleaning robot, is liable and so must compensate injured parties for any harm caused by the operation of the device, even without any breach of a duty of care or any defect with the device.

There would be very few legal defences against this type of strict liability – possibly only a defence of *force majeure* (e.g. if a lorry driven by a third party crashed into the cleaning robot, pushing it against a nearby person who then suffers harm, this might be excluded from strict liability for

¹³ See: Ada Lovelace Institute. (2022). *People, risk and the unique requirements of AI*. Available at: <https://www.adalovelaceinstitute.org/policy-briefing/eu-ai-act/>

the cleaning robot). But most notably, the operator would not escape this strict liability by proving that it has perfectly fulfilled all monitoring, maintenance and similar due diligence obligations.

One problem with this approach is that operators need certainty as to whether they are subject to strict liability or not, which means the scope needs to be clearly defined. Looking more closely at the AI systems that are likely to cause accidents resulting in injury – such as autonomous motor vehicles, aircraft and railways – it becomes clear that most of them are already covered by strict liability under the vast majority of EU legal systems (see scenario 1).

Lawnmowers, cleaning robots and very small drones may in fact be among the few examples where this is not the case, and so the question arises: is it worth creating a whole new regime of AI liability just to deal with these few cases, or would it be simpler to just expand the scope of existing sectoral legislation (such as on vehicles or aircraft)?

Addressing proof-related issues with product liability

A less far-reaching solution would be to tweak product liability regimes of the type introduced by European Council Directive 85/374/EEC¹⁴ in order to improve their application to AI cases. If we look more closely at scenario 2, we realise that the main problem is one of proof. There seems to be a reasonable concern that the person injured by the cleaning or lawnmower robot may not be able to prove it was defective.

However, even if we acknowledge the fact that required standards of proof differ from jurisdiction to jurisdiction, and that there are jurisdictions in Europe where the standard is generally high, the question arises: would a judge not usually rely on doctrines such as *prima facie* evidence? Doesn't the fact that a robot suddenly knocks over a person indicate a defect? Wouldn't the injured party only have to demonstrate that it was the robot which caused the injury by moving towards them (in contrast with a scenario where they themselves stumbled over it?) And is that really so different from traditional accident cases? In other words:

14 European Council. (1985). *Council Directive 85/374/EEC of 25 July 1985 on the approximation of the laws, regulations and administrative provisions of the Member States concerning liability for defective products*. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A01985L0374-19990604>

is it actually likely that a judge would require the injured party to perform a technical analysis of the AI system's code (which would of course be extremely difficult, given the complexity and opacity of machine learning)?

Whether or not the concern is reasonable, explicitly alleviating the burden of proof, or even shifting it to the defendant, could help avoid uncertainty. The future regime for product liability or the AI Liability Directive could explicitly establish that, for example, where an AI-enabled product clearly malfunctions, courts should infer that it was defective and caused the injury.

There could also be certain disclosure obligations with regard to technical specifications and system data. Failure on the part of the defendant to comply with these obligations could then mean a court inferring that the undisclosed information would have been to the defendant's disadvantage.

It ought to be mentioned in this context that Directive 85/374/EEC needs to be revised in any case in order to meet the challenges posed by the digital age (irrespective of AI), for example as its application to standalone software is still uncertain and as it fails to cover defects caused by software updates or their absence only after a product was put into circulation.

Addressing proof-related issues with fault liability

Another policy option would be to tweak fault liability, for example by shifting the burden of proof of fault to the defendant. This could mean requiring that operators of AI systems prove that they did not breach their duty of care. For example, the owner of an autonomous lawnmower or cleaning robot could have to prove that it was produced by a reliable provider, deployed for a task it was designed for, and that adequately trained staff monitored and maintained it according to the instructions.

Once some breach of a duty of care (e.g. omission to install a software update) emerges, alleviations of the burden of proof could also concern the determination of causal links. This could mean, for example, that the defendant would have to prove that the accident would not have been prevented by installing the update. Again, the effects could be further enhanced by way of certain disclosure obligations (concerning technical

specifications and/or logging data), with failure to comply leading to a presumption that disclosure would have been to the defendant's disadvantage.

Such a liability regime would mean very different things for different defendants in different contexts. If it were imposed on consumers using AI-enabled consumer goods the effects could be quite severe, given varying degrees of digital literacy and financial means. For example, this could have the inappropriate result that a consumer operating an autonomous lawnmower on their premises is liable for injury if they were unable to prove that they complied with all instructions, installed all updates, regularly had the device maintained, and so on.

On the other hand, if such a liability regime were imposed on businesses it could work as a welcome additional enforcement mechanism for the AI Act. The company using the cleaning robot in public spaces, for example, would have to meticulously document that it has fully complied with all user obligations resulting from the AI Act (and other legislation) to be able to rebut any presumptions of fault or liability, should injury occur. Of course, any new rules creating a link between the AI Act and liability law could just as well be integrated in the AI Act itself, or in a future product liability regime, as has been illustrated by Chapter III of the European Law Institute Draft for a Revised Product Liability Directive.¹⁵

It should be noted, however, that this would not really close the 'accountability gap' identified above, where operators can 'hide behind the AI', as the operator would still be able to escape liability by demonstrating their due diligence. The effect of this solution would primarily be to enhance compliance with obligations, in particular those resulting from the AI Act, rather than solving problems of under-compensation. Even this effect of enhancing compliance is limited, as it would only show within the range of AI-enabled products that both qualify as 'high risk' within the AI Act and are likely to cause injury to persons. Given that the AI Act is still being negotiated, it is not yet clear how broad or narrow this range of products would be.

15 European Law Institute (ELI). (2022). *ELI Draft of a Revised Product Liability Directive*. Available at: https://europeanlawinstitute.eu/fileadmin/user_upload/p_eli/Publications/ELI_Draft_of_a_Revised_Product_Liability_Directive.pdf

AI liability beyond traditional ‘accident scenarios’

The debate at EU level has focused on ‘accident scenarios’ of personal injury or property damage. By contrast, there are hardly any plans to address the liability issues illustrated by scenario 3 (in the earlier chapter ‘Understanding under-compensation’), i.e. the under-compensation that may result where an AI system causes pure economic loss or non-economic harm. Judging by the questions posed in the public consultation, it seems rather unlikely that the Commission will propose a solution for scenario 3 within the upcoming Directive.

This is all the more surprising as risks of arbitrary decisions, manipulation, discrimination and other similar ‘social risks’ (which the proposed AI Act now calls ‘fundamental rights risks’) are much more specific to the nature of AI systems, compared to the traditional product safety risks related to physical harm, and as this is where the most conspicuous form of ‘accountability gap’ exists

The most plausible explanation for the Commission’s reluctance to address these AI-specific risks is that any potential solution would touch on too many internal spheres of responsibility within the domain of the EU institutions, being relevant across a wide range of areas, from consumer protection to non-discrimination to data protection.

This reluctance is regrettable, as a cross-cutting solution to close the accountability gap would have been fairly straightforward. All that would have been required is a general rule attributing the malfunctioning of AI systems to the person deploying the AI system, in very much the same manner as the malperformance of a human employee is attributed to a corporate defendant. Thus, the bank in scenario 3 would be liable for the credit scoring AI to the same extent as for mistakes committed negligently by one of its clerks.

The message conveyed would have been simple: you cannot escape liability by deploying AI in lieu of employing a human. However regrettably, this solution which was put forward *inter alia* by the Expert

Group on Liability and New Technologies¹⁶ and the German Data Ethics Commission¹⁷ (plus by the author's own response to the public consultation) was not among the possible solutions included in the 2021 consultation and has little chance of being adopted.

16 Expert Group on Liability for New Technologies. (2019). *Liability for Artificial Intelligence*. European Commission. Available at: https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=63199

17 Data Ethics Commission. (2019). *Opinion of the Data Ethics Commission*. German Federal Government. Available at: https://www.bmi.bund.de/SharedDocs/downloads/EN/themen/it-digital-policy/datenethikkommission-abschlussgutachten-lang.pdf;jsessionid=BE6BA606135COD72B4A55BC9566D826A.2_cid322?__blob=publicationFile&v=5

Conclusions and reflections

When discussions about AI and liability started at EU level in 2017, the world was very different from what it is now. Since then, Europe has seen a pandemic, a war, disastrous consequences of the climate crisis, and it is now on the brink of a severe economic crisis. Priorities of policymakers may have shifted over the past five years, and so there may be reluctance to put an additional burden on an EU economy that is already struggling to cope with the current challenges.

Given the current pressures, it would come as a surprise if the Commission decided to adopt the most far-reaching policy option of strict liability. It seems much more probable that we will see a limited and cautious approach being put forward, such as targeted harmonisation across the EU with regard to certain issues, for instance, proof of fault.

This may be a reasonable solution in the given circumstances, but the less added value is provided by the new AI liability regime (as compared with the current situation under national tort laws), and the less the new regime ensures a level playing field for businesses across Europe, the more we may have to discuss whether we need another new piece of legislation at all.

Unless the Commission changes the timeline and postpones publication of the proposal, we will all know for sure on 28 September 2022. Whatever the EU decides, fundamental questions around AI and liability are not going to disappear.

About the author

Christiane Wendehorst is Professor of Civil Law at the University of Vienna as well as Co-Head of the Department of Innovation and Digitalisation in Law. Amongst other functions, she is founding member, Immediate Past President (2017-2021) and since 2022 Scientific Director of the European Law Institute (ELI), Vice-President of the Austrian Jurists' Association (ÖJT) and President of the Division of the Humanities and Social Sciences of the Austrian Academy of Sciences (ÖAW). She is also an elected member of Academia Europea (AE), the International Academy for Comparative Law (IACL), and the American Law Institute (ALI).

At the moment, her research focuses on legal aspects of digitalisation and she has been working as an expert on topics such as digital content, the Internet of Things, AI and data economy, e.g. for the European Commission, the European Parliament, the German Federal Government, the ELI and the ALI. From 2018-2019, she co-chaired the Data Ethics Commission of the German Federal Government.

About the Ada Lovelace Institute

The Ada Lovelace Institute was established by the Nuffield Foundation in early 2018, in collaboration with the Alan Turing Institute, the Royal Society, the British Academy, the Royal Statistical Society, the Wellcome Trust, Luminate, techUK and the Nuffield Council on Bioethics.

The mission of the Ada Lovelace Institute is to ensure that data and AI work for people and society. We believe that a world where data and AI work for people and society is a world in which the opportunities, benefits and privileges generated by data and AI are justly and equitably distributed and experienced.

We recognise the power asymmetries that exist in ethical and legal debates around the development of data-driven technologies, and will represent people in those conversations. We focus not on the types of technologies we want to build, but on the types of societies we want to build.

Through research, policy and practice, we aim to ensure that the transformative power of data and AI is used and harnessed in ways that maximise social wellbeing and put technology at the service of humanity.

We are funded by the Nuffield Foundation, an independent charitable trust with a mission to advance social wellbeing. The Foundation funds research that informs social policy, primarily in education, welfare and justice. It also provides opportunities for young people to develop skills and confidence in STEM and research. In addition to the Ada Lovelace Institute, the Foundation is also the founder and co-funder of the Nuffield Council on Bioethics and the Nuffield Family Justice Observatory.

Find out more:

Website: [Adalovlaceinstitute.org](https://adalovlaceinstitute.org)

Twitter: [@AdaLovelaceInst](https://twitter.com/AdaLovelaceInst)

Email: hello@adalovlaceinstitute.org



Permission to share: This document is published under a creative commons licence: CC-BY-4.0

Preferred citation: Wendehorst, C. (2022)
AI liability in Europe: anticipating the AI Liability Directive. Ada Lovelace Institute. Available at:
<https://www.adalovelaceinstitute.org/report/ai-liability-in-europe/>

ISBN: 978-1-7397950-8-5