ARTICLE

Liability Rules for AI-Related Harm: Law and Economics Lessons for a European Approach

Shu Li[1]* , Michael Faure[2,3] and Katri Havu[1]

[1]University of Helsinki, Faculty of Law, Porthania, Yliopistonkatu 3, Helsinki, 00100, Finland,
[2]Maastricht University, METRO, Bouillonstraat 1–3, Maastricht, 6211 LH, The Netherlands and
[3]Erasmus University Rotterdam, Rotterdam, Zuid-Holland, The Netherlands
*Corresponding author. Email: shuli.eur@outlook.com

Abstract

The potential of artificial intelligence (AI) has grown exponentially in recent years, which not only generates value but also creates risks. AI systems are characterised by their complexity, opacity and autonomy in operation. Now and in the foreseeable future, AI systems will be operating in a manner that is not fully autonomous. This signifies that providing appropriate incentives to the human parties involved is still of great importance in reducing AI-related harm. Therefore, liability rules should be adapted in such a way to provide the relevant parties with incentives to efficiently reduce the social costs of potential accidents. Relying on a law and economics approach, we address the theoretical question of what kind of liability rules should be applied to different parties along the value chain related to AI. In addition, we critically analyse the ongoing policy debates in the European Union, discussing the risk that European policymakers will fail to determine efficient liability rules with regard to different stakeholders.

Keywords: AI-related harm; artificial intelligence; deterrence; developers; law and economics; liability rules; operators; risk-bearing

I. Introduction

Artificial intelligence (AI) not only entails societal benefits but also raises concerns about the risks involved. One of the key questions in this respect is whether current liability regimes have evolved sufficiently to deal with AI-related harm. Most authors agree that, for example, the European Product Liability Directive (PLD)[1] might not suffice for a variety of reasons, the most important one being that the PLD was drafted for addressing mass-produced goods in an era when the manufacturer controlled the risks and was therefore strictly liable for damage resulting from product defects.[2] This idea is outdated from the standpoint of the more dynamic era of new technologies, where not only traditional producers but also software developers and operators may be the cause of product-related harm.

In this article, we discuss the liability for AI-related harm and, in particular, the recent European Union (EU) policy developments in the light of a law and economics framework. The advantage of the law and economics approach is that it provides a clear benchmark

---

[1] 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 [1985] OJ L 210.

[2] See P Machnikowski, "Producers' Liability in the EC Expert Group Report on Liability for AI" (2020) 11(2) Journal of European Tort Law 137, 139.

(social welfare) that helps in choosing suitable liability rules while taking into consideration the most important goals of such rules, namely deterrence and risk spreading. Therefore, our discussion does not touch upon the sphere of distribution issue.

More specifically, we endeavour to address two questions. First, what is, from a law and economics perspective, an optimal allocation of liability for AI-related harm? Second, to what extent do recent EU policy proposals reflect this optimal allocation? Our contribution to the scholarly discussion concerning the regulation of AI is twofold. First, this article provides a law and economics perspective for determining optimal liability rules for different parties along the value chain. Second, based on the economic lessons, we seek to highlight problems and ambiguities within the recent European proposals.

This article contains four substantive sections. We start with a concise discussion concerning the basic questions pertaining to liability for AI-related harm (Section II), and then we present the basic economic model of tort law when dealing with accidents (Section III). We continue our analysis by explaining how this basic economic model can help to determine the most suitable liability rules in the context of AI-related harm (Section IV). Finally, we present a detailed critical analysis regarding the recent EU proposals (Section V).

II. Liability for AI-related harm: setting the scene

1. AI: disruptions and risks

AI refers to a disruptive technology that can achieve specific goals by acting intelligently, in a "human-like" manner. AI differs from traditional pre-programmable software because of its ability to learn, make autonomous decisions and adjust.[3] The notion of an "AI system" is rather broad: if a machine or computer program can act with a degree of autonomy, it can be called an AI system, regardless of whether it is a component of another product or standalone software.

Just as with humans, AI may make errors that cause harm, and this harm may be material or non-material. For example, AI in medical devices or autonomous vehicles (AVs) can physically harm persons or property as well as cause pain or other suffering. AI systems can also cause different economic losses.[4] For instance, a malfunction of a road management system may result in congestion, which can lead to further losses, such as the loss of opportunity of concluding a business contract.

The potential of AI to cause harm calls for an urgent consideration of the ways to reduce such accidents. The first issue in this context is whether or not human beings should be kept "in the loop".[5] More specifically, commentators have raised the question of whether AI systems should be assigned legal personhood in order to hold them accountable for AI-related harm. This approach, however, has been criticised not only because it ignores human oversight, but also due to its infeasibility when one considers, for example, the financial resources needed for compensation.[6] Assigning legal personhood and "personal

---

[3] See, eg, European Commission High-Level Expert Group on Artificial Intelligence, *A Definition of AI: Main Capacities and Disciplines* (2019) p 6; P Čerka, J Grigienė and G Sirbikytė, "Liability for Damages Caused by Artificial Intelligence" (2015) 31(3) Computer Law & Security Review 376; SJ Russell and P Norvig, *Artificial Intelligence: A Modern Approach* (4th edition, London, Pearson 2021) p 19.

[4] Commission, "White Paper on Artificial Intelligence" COM(2020) 65 final, 10–11.

[5] R Koulu, "Human Control over Automation: EU Policy and AI Ethics" (2020) 12 European Journal of Legal Studies 9.

[6] See, eg, IN Cofone, "Servers and Waiters: What Matters in the Law of AI" (2018) 21 Stanford Technology Law Review 167; JS Gordon, "Artificial Moral and Legal Personhood" (2020) 36 AI & Society 457; SM Solaiman, "Legal Personality of Robots, Corporations, Idols and Chimpanzees: A Quest for Legitimacy" (2017) 25(2) Artificial Intelligence and Law 155. Some commentators have also proposed, from a legal standpoint, to treat AI systems as agents of their controllers: R Calo, "Robots as Legal Metaphors" (2016) 30 Harvard Journal of Law & Technology 209.

liability" to AI systems has not received any significant support at the European level.[7] In addition, it is generally agreed that human oversight is necessary and that AI systems should not, at this point, be independent actors from a legal standpoint.[8] The existing policy debates take it as their starting point that regulation should concern the humans and entities that produce and control AI applications. It is therefore imperative to provide these actors with incentives to behave appropriately. Some features of AI, however, make risk allocation a challenging task.

First, many stakeholders and producers are involved, making it a complex task to discern the party that has caused the harm. There are parties that participate in creating a standalone AI system or a product that embeds such a system, but also parties who exercise control over the operation of an AI system.[9] Deciding how the behaviour of each of these parties should be controlled constitutes quite a challenge for regulators.

Second, the decision-making process of an AI system can be opaque. Unlike programmable machines, whose operation is predictable within the limitations of the software, an AI system is a non-transparent "black box".[10] This opaque decision-making may generate more risk that is unknown to stakeholders, and it complicates the assessment of the probability that a harm will occur.[11] Additionally, the number of situations in which every party has behaved appropriately but an accident nonetheless occurs can increase.[12]

Third, the operation of an AI system is fully or partially autonomous.[13] This increasing autonomy blurs the relationship between humans and machines.[14] This may further complicate the determination of optimal liability rules among different parties.

2. Regulating AI-related risks via extra-contractual liability

Reducing risk does not necessarily require legal instruments.[15] Risk can be optimally allocated among different parties through bargaining, if the parties in question are able to bargain with each other in a frictionless manner and all of them hold the perfect information necessary to detect the factors that can affect such risk.[16] In ideal circumstances, any legal intervention is unnecessary and inefficient, because it can result in additional transaction costs.[17]

Legal instruments, however, become indispensable when contracting over risk allocation cannot result in an optimal outcome. If transaction costs are high or where there is information asymmetry, parties might not be able to agree on efficient risk allocation.[18] This happens with AI systems as well. Different parties along the supply chain may represent different legal entities and possess different amounts of information regarding the AI

[7] See, eg, European Expert Group on Liability and New Technologies – New Technologies Formation (EG-NTF), *Report: Liability for Artificial Intelligence and Other Emerging Digital Technologies* (2019) p 37.

[8] See Commission, supra, note 4, 15.

[9] M Buiten, A de Streel and M Peitz, "EU Liability Rules for the Age of Artificial Intelligence" (2021) <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3817520> (last accessed 10 January 2022).

[10] F Pasquale, *The Black Box Society* (Cambridge, MA, Harvard University Press 2015) pp 8–9.

[11] J Burrell, "How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms" 3(1) Big Data & Society 1.

[12] MF Grady, "Unavoidable Accident" (2009) 5(1) Review of Law & Economics 177.

[13] Commission, supra, note 4, 6.

[14] E Talley, "Automatorts: How Should Accident Law Adapt to Autonomous Vehicles? Lessons from Law and Economics" (2019) Hoover Institution Working Group on Intellectual Property, Innovation, and Prosperity <https://www.hoover.org/sites/default/files/ip2-19002-paper.pdf> (last accessed 10 January 2022).

[15] RH Coase, "The Problem of Social Cost" (1960) 3 The Journal of Law & Economics 1.

[16] WY Oi, "The Economics of Product Safety: A Rejoinder" (1974) 5(2) The Bell Journal of Economics and Management Science 689.

[17] GL Priest, *The Rise of Law and Economics: An Intellectual History* (Abingdon-on-Thames, Routledge 2020) p 43.

[18] VP Goldberg, "Toward an Expanded Economic Theory of Contract" 10(1) Journal of Economic Issues 45.

system. Moreover, in some cases victims are third parties that are not in contractual relationships with any of the AI stakeholders. Hence, extra-contractual liability is a crucial tool in providing the parties involved with incentives to behave appropriately, especially when no bargaining can be arranged prior to the accident.[19]

3. Towards European extra-contractual liability rules?

EU-level documents thus far have discussed both liability for AI-related harm in general terms and the more specific question of revising the EU product liability regime. For the purposes of this article, both themes are relevant, as we endeavour to discuss the optimisation of extra-contractual liability rules regardless of whether they are generally applicable or limited to product liability. As of yet, the EU legislator has issued neither general AI liability rules nor final drafts; only tentative plans have emerged.

In 2017, an early Resolution of the European Parliament noted the need to set up a common liability framework in the EU for AI-related harm.[20] A broad Communication on "Artificial Intelligence for Europe" was then introduced by the European Commission in 2018.[21] This resulted, among other things, in the establishment of an Expert Group on Liability and New Technologies. It has two subgroups: the Product Liability Directive Formation and the New Technologies Formation (EG-NTF).

In 2020, the Commission's "White Paper on Artificial Intelligence" stated that the risks involved in the use of AI systems should be addressed by legislative means[22] and underlined the obligations of human actors.[23] Additionally, the Commission Report[24] of the same year emphasised that besides ensuring compensation for victims, liability rules should also "provide *economic incentives* for the culpable parties to avoid causing such damage".[25]

The Final Report of the EG-NTF (2019) sets out that liability rules should cover all key parties along the supply chain as well as the whole lifespan of AI systems. According to that Report, "one reason why existing rules on liability may produce unsatisfactory results is that loss resulting from emerging digital technologies is not allocated to the party *who is the most appropriate* to bear the loss" (emphasis added). It added that losses should be borne by the cheapest cost-avoiders or cheapest insurance-takers.[26] Furthermore, harm caused by a defective emerging digital technology should trigger strict producer liability, irrespective of whether the technology is in tangible or digital form.[27]

In 2020, the European Parliament issued a Resolution (EP Resolution) on civil liability for AI.[28] It called for an evaluation of the PLD[29] as well as setting out a proposal concerning the operators of AI systems. The proposed regime is premised on several key arguments of

---

[19] S Shavell, *Foundations of Economic Analysis of Law* (Cambridge, MA, Harvard University Press 2004) p 175.

[20] European Parliament, "The resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics" (2015/2103(INL)).

[21] Commission, "Artificial Intelligence for Europe" SWD(2018) 137 final.

[22] Commission, supra, note 4, 12.

[23] ibid, 21.

[24] Commission, *Report on the safety and liability implications of Artificial Intelligence, the Internet of Things and robotics*, COM(2020).

[25] ibid, 12 (emphasis added).

[26] EG-NTF, supra, note 7, 5, Key finding [3].

[27] ibid, 42, Key finding [13]. This viewpoint was reiterated by members of EG-NTF in a report drafted by the European Law Institute (ELI). See BA Koch et al, "Response of the European Law Institute: Public Consultation on Civil Liability Adapting Liability Rules to the Digital Age and Artificial Intelligence" (2022) 13(1) Journal of European Tort Law 25.

[28] European Parliament, "Resolution of 20 October 2020 with recommendations to the Commission on a civil liability regime for artificial intelligence (2020/2014(INL))" (2020) <https://www.europarl.europa.eu/doceo/document/A-9-2020-0178_EN.html> (last accessed 10 January 2022).

[29] ibid, para 8.

the EG-NTF Report. A risk-based approach is proposed: operators of high-risk AI systems should be subject to strict liability, and others should be subject to fault-based liability.[30] Turning to the general regulation of AI, in 2021 the Commission published a proposal for the so-called Artificial Intelligence Act (AI Act), which concerns the obligations of those who develop and deploy AI systems.[31] However, this draft AI Act does not contain any comprehensive liability rules.

At present, EU legislators are investigating the necessity of adapting the product liability rules as well as domestic liability rules for AI.[32] The Commission held a public consultation regarding this domain. Over 60% of the received responses indicated that liability rules should be further adapted in order to enhance trust in AI systems.[33]

There has also been an extensive discussion of the liability for AI-related harm in academia. For example, the members of the EG-NTF were invited to further clarify the findings of their report[34] – and they have expressed their views in the abovementioned public consultation.[35] Moreover, the inadequacy of the current product liability rules[36] as well as the problems of the EP Resolution[37] are being expounded by scholars.

III. The basic economic model of determining the optimal liability rule

In general, liability rules are either fault-based or strict. According to the theory of tort law and economics, social welfare serves as the criterion to determine which type of liability rule is appropriate to deal with accidents.[38] One liability rule will be more desirable than the others if it can result in a higher level of social welfare. Social welfare is not only influenced by the deterrent effect of liability rules, but also depends on the risk-shifting effect.[39] In this section, we explain how to identify a desirable liability rule when these two factors are considered.

1. Deterrence as a factor in determining optimal liability rules

From an economic perspective, the primary goal of liability rules is deterrence.[40] In a bilateral accident, where both the tortfeasor and victim contribute to an accident, a socially

---

[30] ibid, Arts 4 and 8.

[31] Commission, "Proposal for a Regulation of the European Parliament and of the Council Laying down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts", COM(2021) 206 final.

[32] See Commission, *Inception Impact Assessment*, Ares (2021)4266516 <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> (last accessed 25 June 2022).

[33] Commission, "Adapting Civil Liability Rules to the Digital Age and Artificial Intelligence: Factual summary report on public consultation" (2022), Ref. Ares(2022)2620305 <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/public-consultation_en> (last accessed 25 June 2022).

[34] See Machnikowski, supra, note 2; C Wendehorst, "Strict Liability for AI and Other Emerging Technologies" (2020) 11(2) Journal of European Tort Law 150; BA Koch, "Liability for Emerging Digital Technologies: An Overview" (2020) 11(2) Journal of European Tort Law 115.

[35] Koch et al, supra, note 27.

[36] See, eg, B Schütte, L Majewski and K Havu, "Damages Liability for Harm Caused by Artificial Intelligence – EU Law in Flux" (2021) Helsinki Legal Studies Research Paper 69 <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3897839> (last accessed 10 January 2022).

[37] A Bertolini and F Episcopo, "The Expert Group's Report on Liability for Artificial Intelligence and Other Emerging Digital Technologies: A Critical Assessment" (2021) 12(3) European Journal of Risk Regulation 1.

[38] S Shavell, *Economic Analysis of Accident Law* (Cambridge, MA, Harvard University Press 2009) p 3.

[39] ibid.

[40] Shavell, supra, note 19, 267–69.

optimal outcome can be achieved if a liability rule can optimise the behaviour of all of the parties.[41] Whether a liability rule can achieve the socially optimal outcome depends on two variables.

On the one hand, social welfare is influenced by the precautionary measures taken by the parties. A desirable liability rule will provide the parties with an incentive to take an optimal *level of care*.[42] This is the point where the marginal costs of a precaution equal the marginal reduction of expected harm. By setting the due care level at this optimal level, either strict liability (with a contributory negligence defence) or a fault-based rule can induce all of the parties to take optimal precautions.[43]

On the other hand, when all of the parties are expected to adopt an optimal level of care within either of the liability rule options, controlling the *activity level* becomes imperative.[44] In this situation, the activity level serves as the key variable to decide which liability rule prevails. Social welfare can be further improved if the relevant parties optimise their level of activity. The factors that have an impact on level of activity (such as personal utility and the intensity and frequency of an activity) can rarely be observed by outsiders.[45] By having the relevant parties internalise the accident cost, liability rules can offer the parties an incentive to optimise their level of activity. Ideally, liability rules incentivise all of the parties to adopt optimal activity level. However, economic models have proved that none of the existing liability rules can induce both parties to optimise their level of activity simultaneously.[46] This is because the residual risk that remains, when all of parties have been diligent, must be allocated to either tortfeasors (under strict liability) or victims (under fault-based liability).[47] Only the parties who are about to bear the residual risk have the incentive to further control their activity level.[48]

In this regard, strict liability and fault-based rules lead to different outcomes when it comes to optimising the level of activity. By applying strict liability to specific tortfeasors, since these parties have to internalise the residual risk given that all parties are diligent, they will have an incentive to further optimise their activity level to minimise the risk borne by them. In contrast, if specific tortfeasors are subject to fault-based liability, they are not liable if they have adopted the required level of care. Hence, under fault-based rules, tortfeasors tend to adopt an excessive level of activity to maximise their personal utility.

From the deterrence perspective, therefore, the value of strict liability is to provide a party with an extra incentive to control their activities, especially when such activities are ultrahazardous with little value addition.[49] However, if adjusting the precautionary level can already appropriately control the behaviour of potential tortfeasors and reduce substantial accident costs, the benefit of controlling the activity level would be marginal, while possibly curbing valuable activities. In this case, fault-based rules would be preferable.

---

[41] Shavell, supra, note 38, 9–10.

[42] S Shavell, "Strict Liability versus Negligence" (1980) 9(1) The Journal of Legal Studies 1.

[43] JP Brown, "Toward an Economic Theory of Liability" (1973) 2(2) The Journal of Legal Studies 323.

[44] Shavell, supra, note 38, 26.

[45] G Dari-Mattiacci, "On the Definitions of Care and Activity Level and the Choice of Liability Rules" (2003) <https://dspace.library.uu.nl/bitstream/1874/723/24/c3.pdf> (last accessed 10 January 2022).

[46] Shavell, supra, note 38, 30–31.

[47] The literature also discusses the possibility of sharing the residual risk between tortfeasors and victims. However, that cannot result in an optimal outcome either. See E Carbonara, A Guerra and F Parisi, "Sharing Residual Liability: The Cheapest Cost Avoider Revisited" (2016) 45(1) The Journal of Legal Studies 173.

[48] G Dari-Mattiacci and F Parisi, "Liability Rules: An Economic Taxonomy" in M Bussani and AJ Sebok (eds), *Comparative Tort Law* (Cheltenham, Edward Elgar Publishing 2021) pp 112–32.

[49] S Shavell, "The Mistaken Restriction of Strict Liability to Uncommon Activities" (2018) 10 Journal of Legal Analysis 1.

2. Risk-bearing as a factor in determining optimal liability rules

The efficiency of a liability rule is not only affected by its deterrent effect, but also by the effect of risk-shifting.[50] Risk-averse parties dislike uncertainties, and they prefer to shift an expected risk prior to accidents rather than to bear the losses after an accident occurs.[51] Social welfare increases if the risk is allocated to a risk-neutral (or less risk-averse) party (if such a party exists) or to a party that is easily able to spread losses. Hence, policymakers need to carefully evaluate the preferences of different parties and possible access to risk-shifting mechanisms.[52]

In reality, there are also additional factors that play a role in risk-shifting.[53] These include (first-party or third-party) insurance as well as risk-sharing agreements. If risk-shifting tools are available, it is less important to allocate the risk initially to the less risk-averse party via liability rules. In other words, the deterrent effect is, then, the primary concern when deciding whether strict liability or negligence-based rules should prevail.[54]

3. How the economic models influenced the liability rules in traditional mass production

In this subsection, we will explain how the liability rules applicable to different tortfeasors have usually been determined in the context of traditional mass production.

In early times, risks related to products arose from manufacturers and consumers. On the one hand, the risk could be created in the process of making the product; on the other hand, the careless behaviour of consumers could result in harm. Consumers and manufacturers could, in any event, easily meet each other in the market, and the former were able to visually evaluate the quality of a product. Therefore, the two parties could de facto allocate the risk by an agreement.[55] This has not been the case since the Industrial Revolution, and currently not only the main manufacturers and consumers, but also other parties in the supply chain could contribute to risks. As a consequence, it is complicated and costly to allocate risks (liability) by utilising contracts.[56] This was the background for introducing specific extra-contractual liability regimes regarding harm related to products.[57] The producers define the main technical features and functions of a product. They are responsible for employing qualified engineers to test the safety of products, examine their technical reliability and store all of the safety proxies and risk reports in the process of production. However, such information may not

---

be equally accessed by courts, regulators and consumers.[58] Considering the substantial risk posed by the activity of producers, providing them with an incentive to optimise their *activity level* is crucial.[59]

In addition, compared to consumers and third parties, producers are more capable of bearing the risk.[60] First, producers are supposed to be less risk-averse than other downstream parties and third-party victims, as they are normally entities (businesses) with assets.[61] Second, having borne the residual risk, without referring to additional risk-shifting mechanisms, producers can further spread the potential losses among consumers via the price mechanism.[62] Considering all of these factors as well as the risk posed by the producers' actions, making producers strictly liable should reduce accident costs and deliver a higher level of social welfare.

Today, as multiple parties are involved in value chains, the situation of subsequent distributors and non-consumer operators must also be considered. These latter parties are responsible for reducing accident costs by controlling their own behaviour. Therefore, liability rules should also provide incentives to these parties along the supply chain. Whereas the main producers can define the essential features of products, these other parties cannot. Most of the risks are posed by their negligence in transiting or using the product. For these reasons, fault-based rules are normally applied to these parties.

IV. Liability for AI-related harm: a theoretical discussion

We will now apply the theoretical conclusions explained above to the AI context. In a value chain related to an AI system, multiple parties can pose a risk either by (contributing to) developing an AI system or by operating such a system.[63] Parties that contribute to the risk should be exposed to liability, taking into account their ability to invest in prevention and their risk-bearing capacity. We will discuss the liability of developers (Section IV.1) and then turn to operators (Section IV.2).

1. The liability of the "developers" of AI systems

Traditionally, it was considered that only producers of physical objects might give rise to risks. This is not, however, true anymore in the digital age, since standalone software, which is not embedded in any product, may pose risks as well.[64] Therefore, the old notion of "producers" – a key concept in the PLD and mainly referring to manufacturers of physical objects – is somewhat outdated.[65] However, the issue here is not whether we should call developers of AI systems "producers". A more fundamental and important theoretical question is whether strict liability or fault-based rules should be applied to all developers.

---

[58] L Kaplow, "Optimal Deterrence, Uninformed Individuals, and Acquiring Information About Whether Acts Are Subject to Sanctions" (1990) 6 Journal of Law, Economics, & Organization 93.

[59] JCP Goldberg and BC Zipursky, "The Easy Case for Products Liability Law: A Response to Professors Polinsky and Shavell" (2010) 123(8) Harvard Law Review 1919.

[60] KN Hylton, "The Law and Economics of Products Liability" 88 Notre Dame Law Review 2457.

[61] GL Priest, "The Invention of Enterprise Liability: A Critical History of the Intellectual Foundations of Modern Tort Law" (1985) 14(3) The Journal of Legal Studies 461.

[62] MG Faure, "Economic Analysis of Product Liability" in P Machnikowski (ed.), *European Product Liability: An Analysis of the State of the Art in the Era of New Technologies* (Cambridge, Intersentia 2016) p 619.

[63] See, eg, S Lohsse, R Schulze and D Staudenmayer, "Liability for Artificial Intelligence" in S Lohsse, R Schulze and D Staudenmayer (eds), *Liability for Artificial Intelligence and the Internet of Things* (Baden-Baden, Nomos Verlagsgesellschaft mbH & Co. KG 2019) pp 16–17.

[64] G Howells, "Protecting Consumer Protection Values in the Fourth Industrial Revolution" (2020) 43(1) Journal of Consumer Policy 145.

[65] See also, eg, Koch et al, supra, note 27, 34.

To address this issue, we need to go back to the criteria that have been explained in Section III.

In the era of AI, digital features have the potential to substantially affect the functioning and performance of a product.[66] The quality of an AI system might even be more important than the physical device.[67] Additionally, the developer of an AI system holds much information that cannot be easily accessed or understood by outsiders.[68] In particular, the utility gained from one algorithmic structure and the risk related to it may only be assessed in an accurate manner by the developer itself. From this perspective, developers should be considered as the party that can avoid the risk connected to the design process at the lowest cost, especially when compared to the victim.[69] Hence, it is imperative to provide developers with an extra incentive to control their activity beyond precautional measures. In addition, developers are often business entities and are therefore actors that have "deep pockets".[70] According to the implications of economic analysis, developers should be placed under an incentive mechanism that is similar to the one applied to traditional producers.[71] By exposing the developers of AI systems to strict liability, a higher level of social welfare is expected to be reached.

In optimising the activity level of developers, the predictability of a risk should be taken into account.[72] Developers' strict liability should be limited to the risk that can be predicted based on the state of the art.[73] These comments bring us back to the point that the primary goal of strict liability is deterrence. If the deterrent effect is marginal, a liability burden that is too heavy will mostly just curb *beneficial* activities.[74]

2. The liability of the "operators" of AI systems

Operators of AI systems are in a position where they can create and reduce risks as well, and this calls for liability rules to provide them with incentives to behave appropriately. A key issue is whether operators should be subject to strict liability or fault-based rules. The features of AI systems as well as the context in which they are deployed complicate the assessment as to which parties should be induced to optimise their activity level. At least three factors should be considered in determining the operators' liability in the era of AI.

a. The severity of risk

When strict liability and fault-based liability are being compared, one crucial question from the deterrence perspective is whether parties should be incentivised to further control their activity level even after the optimal level of precaution has been attained. The answers to this question, however, can vary from one scenario to another, depending on

---

[66] See also G Howells, C Twigg-Flesner and C Willett, "Product Liability and Digital Products" in TE Synodinou et al (eds), *EU Internet Law* (Berlin, Springer 2017).

[67] See, eg, Koch, supra, note 34, 121–22.

[68] Commission, supra, note 24, 6–7.

[69] M Wansley, "The End of Accidents" (2021) 55 U.C. Davis Law Review 269.

[70] W Kowert, "The Foreseeability of Human–Artificial Intelligence Interactions" (2017) 96 Texas Law Review 181.

[71] This does not necessarily mean that the developers of AI systems should be placed under the product liability regime. Whether they are regarded as producers or they will be subject to a different liability regime based on strict liability is a matter of policy choice. The question of institutional design will be touched upon in Section V.

[72] See, eg, MU Scherer, "Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies" (2015) 19 Harvard Journal of Law & Technology 353.

[73] MG Faure, L Visscher and F Weber, "Liability for Unknown Risks – A Law and Economics Perspective" (2016) 7(2) Journal of European Tort Law 198.

[74] See also, eg, A Galasso and H Luo, "Tort Reform and Innovation" (2017) 60(3) The Journal of Law and Economics 385.

the severity of the potential harm related to the activity of an operator. The answer is in the affirmative if the activity of a certain party is ultrahazardous, since an additional incentive of behaving properly can reduce significant social losses.[75] In contrast, if a system only causes other than severe harm, it makes little sense to apply strict liability, which would inhibit useful activities.

It follows that a risk-based approach should be used to determine the liability of operators. In fact, such an approach is nothing novel. For example, legislators have recognised the important roles of those who keep or control dangerous assets such as animals, hazardous substances and vehicles.[76] Strict liability induces these parties to optimise their activity.[77] In most other cases, fault-based liability serves as the default regime for operators.[78]

In the context of AI, however, the question that should be further addressed is: how are we to determine whether a specific application of AI is ultrahazardous or not?[79] The use of AI not only generates conventional safety risks, but also has been linked with, among others, threats to fundamental rights.[80] The law shall further crystallise the scope of ultrahazardous risks with regard to AI. One issue to be noted here is that while the materialised harm of some ultrahazardous risks cannot be easily remedied via the tort system (eg in the form of non-material harm or pure economic losses), it does not influence our evaluation of whether such ultrahazardous risks shall be deterred.[81] Deterrence and compensation are two different issues. We can further develop alternative mechanisms for the purpose of remedy besides the tort regime.

b. The heterogeneity of operators

The heterogeneity of operators raises further issues that should be considered. Traditionally, "operation" means that a party exercises a control over a thing and benefits from owning, deploying or using it. In the era of AI, the parties to which this applies are *frontend-based*. They may generate harm through their decision-making concerning matters such as when and how to operate the AI system, and it is therefore justified to provide the operators with incentives to behave appropriately.[82] As in traditional contexts, operators should be subject to fault-based liability in situations where AI systems cannot be described as particularly dangerous. Strict liability would not produce any significant benefit and might curb valuable activities. Strict liability for frontend operators cannot be easily justified unless the particular AI application is ultrahazardous.

In the context of AI applications, there are a variety of parties that exercise control over AI systems by offering essential and ongoing *backend* services.[83] These make sure that an AI system is able to properly interact with the environment, so that frontend operators as well as third parties can rely on the performance of AI systems. For example, when an AV

---

[75] G Spindler, "User Liability and Strict Liability in the Internet of Things and for Robots" in S Lohsse, R Schulze and D Staudenmayer (eds), *Liability for Artificial Intelligence and the Internet of Things* (Baden-Baden, Nomos Verlagsgesellschaft mbH & Co. KG 2019) pp 125, 140–41.

[76] See also, eg, BA Koch, "The 'European Group on Tort Law' and Its 'Principles of European Tort Law'" (2005) 53(1) The American Journal of Comparative Law 189, 200–01.

[77] See also, eg, G Wagner, "Robot Liability" in S Lohsse, R Schulze and D Staudenmayer (eds), *Liability for Artificial Intelligence and the Internet of Things* (Baden-Baden, Nomos Verlagsgesellschaft mbH & Co. KG 2019) p 50.

[78] G Wagner, "Liability for Artificial Intelligence: A Proposal of the European Parliament" <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3886294> pp 17–18 (last accessed 11 December 2021).

[79] See, eg, Bertolini and Episcopo, supra, note 37.

[80] See, eg, NA Smuha, "Beyond the Individual: Governing AI's Societal Harm" (2021) 10(3) Internet Policy Review 1.

[81] MG Faure and S Li, "Artificial Intelligence and (Compulsory) Insurance" (2022) 13(1) Journal of European Tort Law 1.

[82] See also Bertolini and Episcopo, supra, note 37.

[83] Wendehorst, supra, note 34.

is running on a road, its operation relies on the service provided by various backend parties, such as data providers, cloud service providers, road managers, navigation satellite system providers and so on. They influence the operation of an AI system not only via their strategic behaviour, but also, more importantly, through the functioning of their service.[84] It is difficult to discern whether a certain backend service is included in the same "system" as the actual AI application or whether it is a separable operational process.[85] In any event, the risk can be considered to be posed by the entire interactive system as a whole.

The role of backend operators has gone beyond our traditional understanding of what "operation" signifies, and the "borderline" between "operation" and "creation" is becoming elusive. The practical relevance of backend services in the context of many AI applications suggests that backend operators should be provided with incentives to optimise their activities rather than merely to respect the basic precautionary measures.[86] Considering the way in which developers and backend operators complement each other's activities, it would be strange if they were subject to different liability rules. Unlike frontend operators, backend operators should therefore be subject to strict liability regardless of the nature of the AI application.

c. The level of automation and the degree of control

An AI system may be operated by more than one party. While backend operators will exercise control over the AI system on a continuous basis, frontend operators will determine its actual use. The implications of misbehaviour by frontend operators depend on the degree of control they have over an AI application.[87]

It is usually assumed that by having frontend operators internalise some accident costs (via a fault-based liability) we are able to reduce social costs. This will provide frontend operators with incentives for efficient precaution. This incentive mechanism, however, may not work as expected in the era of AI. As the level of automation increases, the "space left" for a frontend operator to make decisions is limited. If an AI application is highly autonomous, frontend operators are less likely to cause severe accidents by their own actions. Therefore, exposing the frontend operators, in this situation, to a major burden of liability might not really reduce accidents.[88] However, even if a machine is operating autonomously, frontend operators should not be entirely immune from liability. Frontend operators may still contribute to accidents by failures to maintain or update their applications. In other words, a "driver" of an AV will never be a mere passenger considering their duties in safeguarding the operation of the vehicle.[89]

3. Summary

The conclusion based on the law and economics analysis is, with respect to the liability of the developers of AI systems, straightforward: the developers of AI systems can usually be

---

[84] O Dheu, C Ducuing and P Valcke, "The Emperor's New Clothes: A Roadmap for Conceptualizing the New Vehicle" (2020) 75 TRANSIDIT 12.

[85] See also, eg, Spindler, supra, note 75, 127.

[86] See also Koch et al, supra, note 27, 39.

[87] ibid, 26.

[88] See also BA Koch, "Product Liability 2.0 – Mere Update or New Version?" in S Lohsse, R Schulze and D Staudenmayer (eds), *Liability for Artificial Intelligence and the Internet of Things* (Baden-Baden, Nomos Verlagsgesellschaft mbH & Co. KG 2019); see also MF Lohmann, "Liability Issues Concerning Self-Driving Vehicles" (2016) 7(2) European Journal of Risk Regulation 335.

[89] See also JA Pattinson, HB Chen and S Basu, "Legal Issues in Automated Vehicles: Critically Considering the Potential Role of Consent and Interactive Digital Interfaces" (2020) 7(1) Humanities and Social Sciences Communications 1.

considered to have a substantial influence on accident risk. Therefore, it makes sense to hold them (strictly) liable.

The situation is, however, more complex as regards the operators of AI. It is reasonable that operators are liable for accidents, because operators may inflict harm. Nevertheless, all of the operators should not necessarily be subject to the same liability rule. The liability of an operator should be determined by considering three factors: the severity of risk, the heterogeneity of operators and the automation level of an AI system. Taking all of these factors into consideration, frontend operators should, as a starting point, be subject to fault-based liability. However, it is imperative to provide them with an extra incentive (strict liability) to adjust their level of activity where their misbehaviour has the potential to result in significant harm and they exercise a substantial degree of control over the AI system. Therefore, the liability for frontend operators shall be assessed in concrete contexts.

The control exercised by backend operators can be significant. Their actions or omissions can make the entire AI system fail and generate notable social losses. Therefore, it is necessary to provide them with an additional incentive to further control their activity. Backend operators should, as a starting point, be subject to strict liability.

V. Evaluating the European policy developments

In this section, we utilise the theoretical considerations described above in commenting upon policy developments and documents published at the EU level. We will critically analyse first the proposals regarding the liability of developers (Section V.1) and then those regarding the liability of operators (Section V.2).

1. A critical analysis: the liability of developers

As mentioned in Section IV.1, strict liability should be applied to developers considering their essential role in deciding the features of AI systems. This idea is also visible in the EU-level policy discussions.

In general, there are two institutional options to accommodate this incentive mechanism. On the one hand, the already-existing product liability regime can be extended to recognise the developers of AI systems as traditional producers. On the other hand, a separate *sui genesis* liability regime can be designed specifically for AI system developers.[90] The first option has appeared to tempt the legislators. For example, the members of the EG-NTF argued that the PLD should apply to "products regardless of the technology on which they are based".[91] In this case, "developers" of digital goods would be perceived as "producers", and their behaviour would thus be evaluated under the framework of product liability. Strict liability would, therefore, apply to the harm that is caused by *defective* AI systems.

Nonetheless, this approach presents some challenges. According to the existing PLD, only "movables" fall into the scope of "products".[92] Digital goods (such as software) have not been regarded as "products", neither within the meaning of the PLD nor in the national laws of most EU Member States.[93] The limited scope of "product" has precluded software

---

[90] See, eg, Buiten et al, supra, note 9.

[91] Machnikowski, supra, note 2, 139.

[92] Art. 2 of the PLD.

[93] For the attitudes of Member States regarding the issue of whether digital goods are qualified as "products", see Commission, *Evaluation of Council Directive 85/374/EEC on the approximation of laws, regulations and administrative provisions of the Member States concerning liability for defective products: Final Report* (2018) 69–72 <https://op.europa.eu/en/publication-detail/-/publication/d4e3e1f5-526c-11e8-be1d-01aa75ed71a1/language-en> (last accessed 10 January 2022).

developers from being liable under the PLD. Therefore, it should be ensured that harm caused by defective software could be covered under the revised, new product liability rules. Additional challenges, such as how to prove that the harm was *caused* by a *defective* software application, should also be carefully considered.[94]

Another issue to be addressed is whether developers should be responsible for harm if they have adopted an optimal activity level. This question comes close to the issue of whether the so-called *development risk defence* should continue to exist in the AI era.[95] One remarkable change proposed by the EG-NTF is removing the development risk defence from the producers of AI systems.[96] The result would be that a developer is subject to strict liability even if the defect can be detected only after the product is placed on the market.[97] Since an AI system may be subject to some updates, its producer may be considered as the party who can efficiently reduce the risk. It is thus reasonable to have the producer bear some residual risk relating to the obligation to update the AI system in a timely manner. However, the analysis in Section IV.1 has indicated that continuous control only justifies the application of strict liability to developers if the risk can be predicted based on the state of scientific knowledge (which is also developing continuously). With respect to harm that occurs when developers have already optimised their activity, applying strict liability may have a counterproductive effect. It might prevent some accidents, but it could significantly deter beneficial activity as well. Accordingly, the only justification for removing the development risk defence would be that developers have a better risk-shifting capacity. In all, whether the development risk defence shall be retained or not is a matter of compensation more than a question of deterrence. Therefore, a reasonable approach for policymakers would be to pay more attention to developing compensation mechanisms that would be separate from liability rules and applicable to situations where the developers have already optimised their activity.

2. A critical analysis: the liability of operators

This subsection will critically analyse whether the European proposals regarding the liability of operators are efficient. We will analyse first the liability of backend operators (Section V.2.a) and then that of frontend operators (Section V.2.b).

a. Backend operators

As indicated in Section IV.2, the services offered by backend operators can be essential for an AI system. Consequently, a backend service provider often plays a fundamental role. Backend operators should thus be subject to strict liability regardless of their risk profile.

The issue of backend operators has apparently challenged European policymakers. There appear to be divergent views as to whether and how backend operators should be subject to strict or fault-based liability. For example, the EP Resolution proposes that the risk-based approach should be utilised to determine the liability of backend operators. Within this framework, backend operators of high-risk AI systems would be subject to strict liability, while other backend operators would be liable only if a harm is caused by their fault.[98] However, the EP Resolution also indicates that backend operators should

---

[94] See Schütte et al, supra, note 36, 23. There is also literature that suggests that the strict liability of developers should not rely on the notion of "defect". See Spindler, supra, note 75, 136.

[95] Art 7(e) of the PLD. Note that the Member States had the possibility to exclude the defence (Art 15(1)(b)) and therefore the availability of the defence currently varies across the EU.

[96] EG-NTF, supra, note 7, 6, Key finding [14].

[97] ibid, 42.

[98] European Parliament, supra, note 28, Art 4(1).

be covered by the notion of "producers" and be subject to strict product liability regardless of their risk profile.[99]

It seems that the latter approach of coining backend operators as producers is consistent with the lessons from law and economics. This approach, however, is by no means entirely problem-free. For example, the issue of whether backend operators can be equated with producers is complex as such.[100] In addition, strict product liability is nowhere near a "pure" strict liability in the economic sense but is conditioned by additional requirements. Under the European product liability regime, the claimant must in any event prove a "defect".[101] In the future, if policymaking is developed in this direction, it is paramount that this issue of proving a defect is specifically addressed to facilitate claims as well as to avoid legal uncertainty.[102] Otherwise, some social losses might not be internalised by backend operators, and they would, consequently, not be incentivised to behave appropriately. How this issue should be tackled and to what extent it could happen *ex ante* by means of legislation area complex questions and cannot be discussed here in detail.[103]

b. Frontend operators

The severity of risk and the level of automation would, as explained in Sections IV.2.a and IV.2.c, play a central role in determining the liability of frontend operators. Where the activities of frontend operators are hazardous, they should be provided with extra incentives to control their level of activity. In addition, the liability of frontend operators should be calibrated in accordance with their actual degree of control over a specific AI application.

Under the proposal contained in the EP Resolution, frontend operators would be subject to fault-based liability in settings that are not considered high risk. This is consistent with law and economics as well as the overall evolution of the liability of (frontend) operators in other contexts. In the era of AI, policymakers should be careful in defining high-risk AI, also bearing in mind that, in other contexts, "high-risk things" have traditionally been a very small group, including, for example, dangerous chemicals or nuclear power plants.

The EP Resolution does not specify high-risk AI systems, but the newly proposed AI Act offers a wide range of AI applications that would be categorised as high risk.[104] It is uncertain whether a similar list would also be utilised in the context of future liability rules. A list such as the one now published would mean that the scope of strict liability is broad.[105] It is too early to assert whether the use of AI would ultimately extend or reduce the overall scope of ultrahazardous activities recognised by legislation.[106] Some issues should, in any event, be carefully considered and addressed as far as possible. First, for a given accident (eg a traffic accident), victims shall not be treated differently, regardless of whether they

---

[99] ibid, para 8.

[100] See also, eg, Spindler, supra, note 75, 137.

[101] Wagner, supra, note 77, 34–36.

[102] P Fabian et al, "Reasonable, Adequate and Efficient Allocation of Liability Costs for Automated Vehicles: A Case Study of the German Liability and Insurance Framework" (2018) 9(3) European Journal of Risk Regulation 548.

[103] See also, eg, regarding the discussion on rules versus standards, I Ehrlich and RA Posner, "An Economic Analysis of Legal Rule-Making" (1974) 3 Journal of Legal Studies 257; L Kaplow, "Rules versus Standards: An Economic Analysis" (1992) 42(3) Duke Law Journal 557; F Weber, "European Integration Assessed in the Light of the Rules vs. Standards Debate" (2013) 35(2) European Journal of Law and Economics 187.

[104] Commission, "Proposal for a Regulation of the European Parliament and of the Council Laying down Harmonised Rules on Artificial Intelligence (ARTIFICIAL INTELLIGENCE ACT) and Amending Certain Union Legislative Acts", COM(2021) 206 final, Annexes 2 & 3.

[105] For example, while AVs could reduce some risks caused by the tiredness of drivers, they may generate risks that drivers could easily avoid. See Wagner, supra, note 77, 44.

[106] ibid, 46.

have been harmed by AI or not.[107] This in nature is a compensation issue. From a deterrence perspective, the equal treatment of victims does not mean that the liability rules for all tortfeasors should be identical. In the era of AI, considering the different automation levels of an AI system in a given scenario, the degree of control exercised by frontend operators is correspondingly different. Therefore, it is reasonable to crystallise liability rules in line with the degree of control. Second, the existing European strict liability rules are mostly rules of national laws.[108] While harmonising the liability rules for frontend operators of AI systems may reduce the divergence between the EU Member States concerning this specific matter, the novel rules and pre-existing national liability rules on dangerous activities might form an incoherent whole within each Member State,[109] even if a highly fragmented legal landscape and uncertainty about applicable rules can also be problematic and costly.

Moving on, our next question is whether the European proposals take the automation level into account appropriately when determining the liability of frontend operators. The EP Resolution noted that, where there are multiple operators, "the proportion of liability shall be based on the respective degree of control the operators had over the risk connected with the operation and functioning of the AI system". Additionally, the EG-NTF explained that "strict liability should lie with the person who is in control of the risk connected with the operation of emerging digital technologies", and if there are multiple operators, "strict liability should lie with the one who has more control over the risks of the operation".[110] This language is problematic, since backend operators are always the parties that can efficiently reduce social losses caused by backend services, regardless of the degree of control exercised by frontend operators. It seems that both the EP Resolution and the EG-NTF Report are determined to find the *cheapest cost-avoider* between frontend and backend operators, even if these two parties are not comparable with each other and cannot be considered mutually exclusive.

From an institutional perspective, what really matters is determining the (burden of) liability of frontend operators in accordance with their degree of control over an AI system. Regarding this issue, there is no "one-size-fits-all" solution for all AI applications.[111] The liability of frontend operators can only be accurately decided with an approach that is both sector-based and evidence-based.[112] Because of this, merely adopting AI technology does not lead to any clear, automatic implications in terms of altering the existing liability rules.

VI. Conclusion

In this article, we explained how a law and economics analysis can contribute meaningfully to the debate about the liability rules that should be applied to AI-related harm. In the EU,

---

[107] Wendehorst, supra, note 34, 173.

[108] Wagner, supra, note 77, 49.

[109] Lohsse et al, supra, note 63, 18; see also Wagner, supra, note 78, 18–19.

[110] EG-NTF, supra, note 7, 39, Key finding [11].

[111] E Karner, "Liability for Robotics: Current Rules, Challenges, and the Need for Innovative Concepts" in S Lohsse, R Schulze and D Staudenmayer (eds), *Liability for Artificial Intelligence and the Internet of Things* (Baden-Baden, Nomos Verlagsgesellschaft mbH & Co. KG 2019) pp 117, 122; see also Bertolini and Episcopo, supra, note 79.

[112] For example, according to SAE International, all on-road motor vehicles are categorised into six levels in accordance with their level of automation. Correspondingly, the degree of control exercised by frontend operators is also different. The liability of frontend operators must be adjusted to the degree of control. The extent to which a frontend operator is subject to which liability, however, can only be decided in a sector-by-sector manner. Soft laws, such as standards, can help crystallise the concrete role taken by a frontend operator. See SAE, J3016 Visual Chart <https://www.sae.org/blog/sae-j3016-update> (last accessed 10 January 2022).

we have seen several policy initiatives and preparatory documents concerning such harm. Notably, there has not been any clear, final set of future rules on the table yet. The situation remains open in terms of what kind of liability rules – if any – will be enacted in the near future by EU legislators.

There are several comments that can nevertheless be put forward in terms of ideas and discussions observable in EU-level policy papers. Most importantly, the official proposals have not been entirely clear and consistent on the question of who exactly should be strictly liable and in which particular situations.[113] However, this question is key to the formulation of liability rules. A welcome sign, in any event, is that, for instance, the Commission documents that have set out preliminary ideas as to liability seem to recognise that the efficiency perspective is crucial when deciding what kind of liability framework should be adopted (or retained).[114]

Our analysis has underlined that *developers* exert a considerable influence on accident risk and often also have "deep pockets". Accordingly, they should be subject to strict liability. Such a regime would generate incentives for both optimal care and optimal activity levels as well as facilitating the effective spreading of risk.

Similarly to other frontend operators, the frontend operators of AI systems should, in principle, be subject to fault-based liability, unless their activities could cause significant social losses. The risk-based approach proposed by EU legislators is a step in the right direction,[115] but it should be further developed to clearly define the situations in which frontend operators are strictly liable. This matter should be addressed on a sectoral basis. It is not possible to establish a harmonised, one-size-fits-all liability regime regarding this question at the EU level. Additionally, it is important to assess the current Member State laws on civil liability and evaluate whether they lead to satisfactory results, as well as whether there would be too much divergence across the Union from the standpoint of the single market.[116]

Moreover, the issue of the liability of backend operators should be considered separately from the position of frontend operators. The essential role played by backend operators signifies that they are a party "creating" rather than "operating" AI systems. Strict liability would ensure that backend operators have appropriate incentives.

Several further comments can be made. For example, current EU policy documents focus on identifying "the appropriate party"[117] that should bear risks. However, from an economic perspective, there need not be a single appropriate party – all actors that influence accident risk should be incentivised to take preventative measures. The control exercised by frontend operators does not eliminate the need to provide incentives to backend operators. The relevant question is whether all parties involved should be exposed to strict or fault-based liability and whether the liability of various stakeholders should be separate or rather solidary. In addition, problems could arise in the case of the insolvency of one of the stakeholders. For this reason, the EU documents published also propose the introduction of compulsory solvency guarantees such as mandatory liability insurance.[118] Whether the compulsory purchase of liability insurance for relatively new risks, such as AI-related harm, is justified is open to debate.[119] It is also noteworthy that seeking maximum deterrence may well open the compensatory floodgates. Compensating AI-related damage by other means and mechanisms than by mere liability rules merits further

---

[113] See, eg, European Parliament, supra, note 28; Commission, supra, note 24.
[114] See Commission, supra, note 24, 12.
[115] See, eg, European Parliament, supra, note 28.
[116] See also, eg, Koch et al, supra, note 27.
[117] EG-NTF, supra, note 7, 39, Key finding [11].
[118] European Parliament, supra, note 28, Art 4(4).
[119] Faure and Li, supra, note 81.

attention. These questions as well as many others would undoubtedly benefit from further research.

Competing interests. The authors declare none.