

AI regulation as corporate regulation: accounting for irresponsibility^a

*Michelle Worthington**

Abstract

While questions of responsibility, including legal responsibility, inevitably arise in the process of designing effective AI regulation, it is considerations of irresponsibility, and corporate irresponsibility in particular, that offer regulators the clearest insights into regulatory possibilities. In this article I argue that the process of designing AI regulation (including the necessary process of allocating legal responsibility for AI related harms), is best approached as a subset of corporate regulation, where anticipating and guarding against corporate irresponsibility is a key function of the regulator. Unless it is properly sensitised to the design of corporate legal personality, especially that of the Anglo-American style corporation, AI regulation will be vulnerable to being obstructed by the operation of an intrinsic and irresponsible distributive function sitting at the heart of the corporate device.

Keywords: artificial intelligence (AI), regulation, risk-based regulation, corporations, responsibility, corporate irresponsibility.

Sebbene nel processo di elaborazione di una regolamentazione efficace dell'intelligenza artificiale sia imprescindibile affrontare le questioni attinenti alla responsabilità, inclusa quella giuridica, sono le considerazioni legate all'irresponsabilità – e in particolare all'irresponsabilità d'impresa – a offrire ai regolatori le indicazioni più chiare sulle potenzialità normative. In questo contributo si sostiene che la progettazione della regolamentazione dell'IA (compreso il necessario processo di attribuzione della responsabilità giuridica per i danni connessi all'utilizzo dell'intelligenza artificiale) debba essere intesa come un sottoinsieme della regolamentazione societaria, in cui la previsione e la prevenzione di condotte irresponsabili da parte delle imprese rappresentano una funzione centrale dell'attività regolatoria. In mancanza di un'adeguata comprensione della struttura giuridica della personalità societaria – in particolare nel modello della corporation anglo-americana – la regolamentazione dell'IA rischia di essere compromessa dal funzionamento di una funzione distributiva intrinsecamente irresponsabile, insita nel nucleo stesso del modello societario¹.

^a Received 29/07/2025 and published 09/12/2025.

* Australian National University Law School, e-mail: michelle.worthington@anu.edu.au.

¹ This abstract was translated from English into Italian using ChatGPT4 (July 2025 version).

Parole chiave: intelligenza artificiale (IA), regolamentazione, regolamentazione basata sul rischio, imprese, responsabilità, irresponsabilità d'impresa.

1. Introduction

Since the start of the decade securing appropriate artificial intelligence (AI) regulation has emerged as a priority regulatory issue in every major jurisdiction around the globe. In no small part, the urgency surrounding this new regulatory challenge is linked to a perception that a hierarchy of influence is establishing itself. In particular, there is a sense that regulatory postures adopted at this juncture will help to determine whether a given jurisdiction becomes a 'maker' or 'taker' of AI regulation², with attendant positioning within the developing AI global order.

For the governments, intergovernmental bodies and other regulatory agencies charged with deciding whether and/or how to regulate AI, issues of *responsibility* loom large. In one sense this is entirely unremarkable. Much of the law in any jurisdiction concerns the allocation of responsibility, either in the form of *prospective legal responsibility* or *historic legal responsibility*³. That responsibility would become a key regulatory concern vis-à-vis AI regulation is to be expected. In fact, it is unavoidable. When viewed from a high level of abstraction the process of devising regulatory schemes aimed at emergent AI technologies can be understood broadly as a process of establishing new *AI responsibility practices*⁴, and legal responsibility practices in particular. This is particularly true of anticipated AI *harms*. Establishing responsibility practices for AI harms, including proper distributions of risk, is a significant part of the regulatory task.

But while regulators are used to working closely with notions of responsibility, it would be a mistake to view the task of regulating AI as a familiar regulatory task. The questions around responsibility that arise in the context of AI technology are meaningfully distinct from responsibility-relevant questions that present in other regulatory contexts. This is for the simple fact that AI tends to confound existing conceptions of responsibility in both morality and law⁵.

The responsibility-confounding qualities of AI are well documented. For present purposes it is sufficient to note, following Chesterman, that the combination of 'speed'⁶, 'autonomy'⁷ and 'opacity'⁸ exhibited by various AI systems greatly challenges existing moral and legal structures, not least as this concurrence of qualities complicates the assigning of responsibility⁹, including legal responsibility, for AI related harms. The European

² UK Government, *AI Opportunities Action Plan*, January 13, 2025, p. 6.

³ P. Cane, *Responsibility in Law and Morality*, Hart Publishing, Portland Or. 2002, pp. 32-33.

⁴ Ivi, pp. 4, 56-60.

⁵ See generally F. Santoni de Sio & G. Mecacci, *Four Responsibility Gaps with Artificial Intelligence: Why they Matter and How to Address them*, «Philosophy and Technology», 34, 2021, pp. 1057-1084.

⁶ S. Chesterman, *We the Robots: Regulating Artificial Intelligence and the Limits of the Law*, Cambridge University Press, Cambridge 2021, chapter 1.

⁷ Ivi, chapter 2.

⁸ Ivi, chapter 3.

⁹ See generally Ivi, chapter 4.

Commission's decision in February to withdraw the EU AI Liability Directive¹⁰ illustrates both the scope of the challenge that regulators currently face, and the scope of disagreement that persists with respect to how particular polities should approach questions of responsibility for AI related harms.

In what follows I argue that the process of designing AI regulation (including the process of allocating legal responsibility for AI related harms), is perhaps best approached *indirectly*. More particularly, I argue that if AI technology confounds conceptions of responsibility in morality and law, then the obverse concept – *irresponsibility* – may prove a more fruitful lens through which to view questions of AI regulatory design. This is for two reasons.

1) First, in its current state of development, AI technology may be regarded as irresponsible in a radical sense. Burgeoning hype around 'agentic'¹¹ AI notwithstanding, AI technology remains instrumental in character. It is relevantly *irresponsive* and *unresponsive*¹², a tool incapable of experience and incapable of choice, at least in the manner that is native to conventional understandings of moral personhood¹³. Ascribing notions of moral responsibility to [current] AI systems is a 'category mistake'¹⁴.

2) Second, the context in which AI systems are being developed and deployed is overwhelmingly corporate in its complexion. It is, therefore, a context characterised by irresponsibility, at least insofar as that concept is understood here. Indeed, some forms of corporation – such as the Anglo-American style for-profit corporation – may be understood to be irresponsible *by design*¹⁵. As will be discussed below, in that form of corporation irresponsibility can be seen to emerge from the ways in which the corporate device distributes the risks and rewards associated with corporate conduct¹⁶. Rewards are largely distributed to those 'internal' to the corporation, whereas risks (including risks of harm) are largely distributed amongst those 'external' to the corporation¹⁷. What we can think of as the corporation's irresponsible distributive function has important implications for AI regulation, particularly AI regulation that is aimed at controlling for the risks associated with AI.

¹⁰ A. Datta & T. Hartmann, *Commission withdraws AI liability directive after Vance attack on regulation*, «Euractiv», February 11, 2025, <https://www.euractiv.com/section/tech/news/commission-withdraws-ai-liability-directive-after-vance-attack-on-regulation/>.

¹¹ See for example T. Finn & A. Downie, *Agentic AI vs. generative AI*, «IBM Think», <<https://www.ibm.com/think/topics/agentic-ai-vs-generative-ai>>.

¹² E.g. J. Gardner, "Relations of Responsibility" in *Crime, Punishment and Responsibility: The Jurisprudence of Antony Duff*, R. Cruft (ed) Oxford University Press, Oxford, 2011, p. 87; See P. Monti, *AI Enters Public Discourse: a Habermasian Assessment of the Moral Status of Large Language Models*, «Ethics & Politics», 61, 1, 2024, pp. 61-80.

¹³ Moral personhood being broadly recognised as a pre-condition for ascriptions of moral responsibility. For a helpful treatment see F. Rudy-Hiller, *The Epistemic Condition for Moral Responsibility*, in «The Stanford Encyclopedia of Philosophy» (Winter 2022 ed), E.N. Zalta & U. Nodelman (eds.), <https://plato.stanford.edu/archives/win2022/entries/moral-responsibility-epistemic/>.

¹⁴ N.H. Conradie and S.K. Nagel, *No Agent in the Machine: Being Responsible and Trustworthy about AI*, «Philosophy and Technology», 37, article 72, 2024, p. 5. Causal responsibility, of course, may be attributed to AI.

¹⁵ E.g. P. Ireland, *Limited Liability, Shareholder Rights and the Problem of Corporate Irresponsibility*, «Cambridge Journal of Economics», 34, 5, 2010, pp. 837-856.

¹⁶ K. Greenfield, *Saving the World with Corporate Law*, Boston College Law School Research Paper 130, 2007, pp. 10-11, <http://dx.doi.org/10.2139/ssrn.978242>.

¹⁷ *Ibidem*.

This piece explores aspects of the second claim. This treatment is intended to signal – but not exhaustively establish or explain – potential regulatory obstacles. In what follows I argue AI regulation must accommodate corporate irresponsibility, including the inherent, irresponsible distributive function exhibited by certain forms of the corporate device. Properly understood, AI regulation is a sub-set of corporate regulation. With few exceptions, the most appropriate point of departure in discussions around the design of AI regulation is not ‘algorithms’, but ‘corporations’, together with everything we have learned about corporate irresponsibility. The task of devising regulatory schemes for AI should be approached accordingly.

2. Some definitions

For the purposes of this discussion the terms ‘regulation’, ‘regulatory schemes’ and ‘regulatory frameworks’ means ‘public control of a set of activities’¹⁸. Legislation (including subordinate legislation) would be the preeminent example of regulation under this definition¹⁹. The concept of ‘AI regulation’ should be read consistently with the above as formal rules devised and enforced by government concerning the development and/or use of AI technology. AI regulation need not take the form of horizontal or general regulation – it can operate in relation to discrete areas of the law. Existing regulation that does not specifically target AI systems but nonetheless captures such systems will not be AI regulation.

The concepts of regulation and AI regulation discussed in this piece are related to, but distinct from, the concept of the ‘regulatory task’, a concept discussed in Part 3.

‘Artificial intelligence’ means a “machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions...[they] vary in their levels of autonomy and adaptiveness”²⁰.

‘Responsibility’ is a word that houses many meanings²¹ and here a number are invoked. In particular the discussion in this paper concerns notions of *legal* responsibility, both prospective and historic. Prospective legal responsibility exists if one has a legal duty to engage in certain conduct, or to avoid engaging in certain conduct²². Prospective legal responsibility has a forward looking, normative function²³. Historic legal responsibility is backwards looking, is evaluative in nature, and concerns legal rules that ascribe legal accountability and/or liability for certain conduct, including sanctions²⁴. In Part 6, the

¹⁸ S. Chesterman, *We the Robots*, cit., p. 4.

¹⁹ ‘Soft law’ as a form of private regulation falls outside the scope of ‘regulation’ in this piece.

²⁰ OECD, *Explanatory Memorandum on the Updated OECD Definition of an AI System*, «OECD AI Papers», 8, 2024, p. 4.

²¹ For a helpful treatment of the emergence and trajectory of responsibility as a philosophical concept in western culture see generally R. McKeon, *The Development and the Significance of the Concept of Responsibility*, «Revue Internationale de Philosophie», 11, 39(1), 1957, pp. 3-32.

²² P. Cane, *Responsibility in Law and Morality*, cit., pp. 31-32.

²³ Ivi, p. 57.

²⁴ *Ibidem*.

discussion also invokes a conception of moral responsibility taken from Williams, where responsibility ‘represents the readiness to respond to a plurality of normative demands’²⁵.

The term ‘distributive function’ is used here in two principal ways: first to refer to the distribution of risk and reward (between different sections of the community) that is produced by the operation of particular legislative or regulatory instruments²⁶; and second, to refer to the distribution of risk and reward that is produced by the operation/design of corporate legal personality. As will be discussed in parts 5 and 6, particular forms of legal personality, such as corporate legal personality, are imbued with their own, internal programming around the distribution of risk and reward. As mentioned above, in the context of the Anglo-American style for-profit corporation, that distributive function operates predictably to internalise reward and externalise risk. The interaction between these two forms of distributive function (i.e. as it emerges from legislation/regulation, and as it emerges from corporate legal personality) is a central concern in the discussion below.

3. *The regulatory task*

Analytical engagement with regulatory phenomena is necessarily built upon a particular understanding or conception of what can be thought of as the relevant ‘regulatory task’²⁷. The regulatory task that is the focus of this paper concerns governmental exercise of its regulatory function vis-à-vis emerging AI technologies in governance contexts characterised by representative government and market-based economic structures. By necessity, the regulatory task includes negotiation/debate informing the design and implementation of regulatory schemes. It may also include preliminary or ‘placeholder’ steps taken by governments that stop short of regulation, but which are nonetheless intended to influence the development and use of AI technology (codes of conduct, principles, white papers, etc).

Within such governance contexts, there appears to be two overarching concerns animating the regulatory task. These are respectively: 1) promoting innovation in AI; and 2) ensuring there are rules around the development and use of AI to mitigate against, or provide redress for, harms associated with the technology²⁸. Given the nature of the technology and its possible applications, there is considerable tension between these two concerns. Indeed, at least insofar as innovation is understood to mean the development of novel technologies, then these two concerns may be regarded as largely antagonistic. While the scope of the regulatory task will evolve over time, in this early phase the task appears to involve identifying and establishing the legal landscape in which both of the above concerns may be pursued, albeit perhaps asymmetrically.

So conceived, and viewed from a high level of abstraction, the regulatory task is not limited to the striking of a balance between individual freedom and public interest (as is so often the case, particularly in contexts involving representative government). Rather, the

²⁵ G. Williams, *Responsibility as a Virtue*, «Ethical Theory and Moral Practice», 11, 4, 2008, pp. 455–470, p. 459.

²⁶ See for example P. Cane, *Responsibility in Law and Morality*, cit., p. 186-190, 219.

²⁷ See for example J. Black & R. Baldwin, *Really Responsive Risk-Based Regulation*, «Law & Policy», 32, 2, 2010, pp. 204-205.

²⁸ E.g. UK Government, *A Pro-Innovation Approach to AI Regulation*, (White Paper) 2023, <<https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach/white-paper>>; EU AI Act, recitals 1 and 2.

regulatory task is one that Cane might describe as ‘complex’²⁹ - it involves balancing *some* individual freedom against public interest, but also it involves attempting to balance *different, potentially competing public interest gains*. Here I am referring to the (largely anticipated) public interest associated with promised technological advancements in AI, as well as the obvious public interest that vests in preserving the safety of individuals, of communities, and the systems that both are reliant upon.

Understood in this way, the regulatory task is one that is strongly *distributive in nature*³⁰. Which is to say, the task facing regulators is one in which they will take a hand in distributing the likely risks and likely benefits of AI technologies between different sections of the regulated community³¹.

4. Distribution of risk and reward

In regulatory contexts involving market-based economies, the distribution of reward is one that is to a great extent shared with ‘the market’. For example, the creation and provision of AI products/services, as well as the aggregated evaluations and choices of consumers with respect to those things, comprises a systemic crucible from which reward emerges. Likewise, the ultimate allocation of profits generated from AI market activity is ordinarily a matter of private ordering. Active governmental involvement in the distribution of reward will typically take the form of taxation, as well as legal protections for workers.

The fact that in market-based economies governments are content to share the task of distributing reward may help to explain why positive steps taken in pursuit of the regulatory task are largely concerned with distributions of risk. Notably, innovation friendly, harm mitigation as opposed to elimination³² style regulatory postures (harm mitigation approaches) have already proven popular, and are perhaps the paradigmatic example of this early trend³³.

In a number of key jurisdictions, regulators appear to be framing the regulatory task with respect to AI as one that largely involves: first, identifying possible risks of harm associated with the development and use of AI (including economic, social and environment harms etc), as well as risks of harms associated with poorly designed regulation of such technology; second, deciding precisely how those identified risks of harm are to be distributed within a particular community; and third, designing and establishing legal responsibilities (both prospective and historic) aimed at producing the desired distribution of risk.

The above description highlights the central role played by considerations of responsibility in the design and implementation of AI regulation. The third, functionally distributive step described above involves regulators settling upon specific forms of *legal* responsibility so as to give expression to *moral/normative* allocations of responsibility agreed upon in the second step (with a particular focus there on questions around moral agency, such as accountability for harmful conduct, and moral obligation to avoid harmful conduct

²⁹ P. Cane, *Responsibility in Law and Morality*, cit., p. 251.

³⁰ Ivi, pp. 282-283.

³¹ P. Cane, *Responsibility in Law and Morality*, cit., p. 186-190, 219; S. Chesterman, *We the Robots*, cit., 177-180.

³² C. Coglianese, *The Law and Economics of Risk Regulation*, University of Pennsylvania, Institute for Law & Economics Research Paper No. 20-18, 2020, p. 9.

³³ Harm mitigation approaches are emerging in a number of jurisdictions, including the EU, Australia the UK, Canada and Singapore.

where feasible). Such a pattern of engagement with considerations of moral and legal responsibility is one that is deeply familiar in the regulatory context, especially in circumstances involving regulation intended to have a strongly distributive function.

5. Regulatory indifference on issues of legal personality

A final observation that can be made about the regulatory task is that regulators in various jurisdictions do not appear to be particularly concerned with sensitising performance of the regulatory task to the *identity* of the subjects of regulatory intervention. More particularly, regulators appear to be approaching the regulatory task armed with the tacit assumption that legal persons will respond to legal responsibilities created by AI regulation in a consistent manner.

Consider for example the EU's *AI Act*. This Act imposes various legal responsibilities upon 'providers', 'deployers', 'importers', 'product manufacturers' etc of AI technology³⁴. Under Article 3 of the Act a 'provider', 'deployer', 'importer' etc could be a natural human being, or alternatively, various different forms of artificial/institutional legal personality such as for profit corporations, or public authorities. The default design of the law is to function in an undifferentiated way in relation to legal personality, absent some specified carve-out or exception. This approach arguably evidences a belief that the true levers of regulation are properly crafted legal responsibilities, where successful performance of the regulatory task involves getting the 'right' mix of prospective and historic legal responsibilities to target the 'right' legal persons.

This view of the regulatory task is not incorrect, but it is incomplete. Legal persons are *not* functionally neutral units. It is simply *not* the case that the imposition the same legal responsibilities across a range of different legal personalities can be expected to produce a desired behavioural result in a consistent manner. Different categories of legal personality – natural legal personality, corporate legal personality and bodies politic – typically exhibit quite distinct legal designs. Relevant differences include the purposes for which the grant of legal personality is made, the various legal capacities and obligations (including perhaps responsibilities) comprising the legal person, and finally, the existence of certain specified conditions imposed against a grant of legal personality, including conditions for use³⁵. As might be expected, these differences in design produce differences in behaviour.

Critically for present purposes, a category of legal personality can come with its own internal or intrinsic programming with respect to *responsibility*, both moral and legal. This programming around responsibility can in turn create an intrinsic *distributive function* within a given category of legal person. If it is not taken into account, this intrinsic distributive function can obstruct the operation of external regulation where such regulation is intended to have a distributive effect. In Part 6 below I signal how the distributive function inherent in the Anglo-American style for-profit corporation – a distributive function which may be described as deeply *irresponsible* – may impede aspects of the regulatory task.

Focussing on this particular style of corporation is a necessary function of context. Unlike the industrial revolution in which unincorporated business models such as

³⁴ See for example *EU AI Act* article 2(1). These legal responsibilities are largely prospective in nature.

³⁵ M. Worthington & P. Spender, *Constructing Legal Personhood: Corporate Law's Legacy*, «Griffith Law Review», 30, 3, 2021, pp. 348-373.

partnerships played an important role³⁶, the AI revolution is predominately driven by large, even global, corporate interests. With the United States having emerged as an early leader in the development of AI, for-profit Anglo-American style corporation is playing an outsized role in shaping the nature and use of AI technologies.

6. *Complicating the regulatory task – corporate legal personality and corporate irresponsibility*

As noted above, understanding the intrinsic distributive function of the for-profit Anglo-American style corporation (from hereon in ‘the corporation’) is a matter of tracing the pattern of legal responsibilities – both prospective and historic – embedded within the design of the device. The configuration of these legal responsibilities in turn reflects determinations about the pattern of moral responsibility accorded to the corporation. With respect to the corporation, the most striking feature of each of these patterns of responsibility, legal and moral, is the extent to which they are oriented around a single concern – corporate financial self-interest.

Tracing first the pattern of legal responsibility, or irresponsibility, as it emerges from the corporation, there are four overlapping elements to consider.

First, corporate decision-making is subject to the operation of key prospective legal responsibilities, including the legal duty imposed upon directors that for convenience I will call the Best Interests Duty (BID). The precise content of the BID is jurisdiction dependent, however it is reasonable to describe it as requiring directors to act always in the best interests of the corporation, where corporate interests are ordinarily understood to equate to, or correlate with, the financial interests of the corporation’s shareholders³⁷. Critically, as shareholder interests are served by the creation of corporate profit, the BID operates to ensure that profit generation is prioritised above *all other considerations* in corporate decision-making. It is difficult to overstate the extent to which the BID trains directorial attention towards corporate financial self-interest. As noted by Greenfield:

managers are held – or consider themselves held – to an obligation to take care of shareholders even when it hurts other stakeholders or society at large, and even when the benefits to those shareholders do not outweigh the costs to others³⁸.

It is even possible, though the question remains subject to debate, that under the BID ‘managers have an obligation to violate *external* laws when necessary to meet their *internal* obligations to maximize returns to shareholders’³⁹. As profit generation is aided by reducing the corporation’s expenditures, the BID ensures that where possible, the corporation will *externalise its costs*, visiting them upon employees and their families, consumers, civic and political institutions, society at large (both present and future), other species, and the natural environmental systems upon which we all depend.

³⁶ P. Ireland, *Limited Liability, Shareholder Rights and Corporate Irresponsibility*, cit., p. 839.

³⁷ E.g. s172 *Companies Act 2006* (UK); s181 *Corporations Act 2001* (Cth); *Revlon Inc. v MacAndrews & Forbes Holdings Inc* 506 A.2d 173 (Del 1986).

³⁸ K Greenfield, *Saving the World with Corporate Law*, cit., p. 12.

³⁹ Ivi, p. 12, citing F.H. Easterbrook & D. R. Fischel, *Antitrust Suits by Targets of Tender Offers*, «Michigan Law Review», 80, 1982, pp. 1155-1178, p. 1177, n.57.

The operation of the BID is reinforced by the nature and function of shareholding in the corporation. This brings us to the remaining three elements for consideration, limited liability, separate legal personality and control rights.

Under the doctrine of limited liability, a shareholder's financial liability is limited to the cost of their investment in the corporation, i.e. the full price of their shares. Critically, a shareholders' personal assets are generally protected from claims relating to corporate liability. Moreover, separate corporate legal status ensures that it is the corporation, and not its shareholders, that has historic legal responsibility for corporate debts/liabilities. In the face of corporate collapse or corporate malfeasance, shareholders can lose the value of their investment, but nothing more. So shielded from legal responsibility, shareholders are far less likely to be concerned with the impacts, including negative impacts, of corporate conduct⁴⁰.

The nature and structure of corporate control rights compounds this phenomenon. As legal responsibility for day-to-day corporate management is ordinarily reserved for the board of directors, shareholders continue to derive an income stream irrespective of whether they have formed an understanding as to how corporate profits are generated⁴¹. Despite this fact, they also retain ultimate control over the corporation's activities – directors are accountable to shareholders, not least shareholders are (as a general proposition) legally responsible for the appointment and removal of directors. Moreover, shareholders are empowered to hold directors to their directorial obligations, including their obligations under the BID; an internal enforcement mechanism that helps to ensure that the corporation maintains its single-minded focus on profit generation⁴².

How then, might we describe the pattern of moral responsibility that is revealed by the design of legal responsibility in the corporation? If the concept of responsibility is understood in the manner suggested by Williams, where responsibility 'represents the readiness to respond to a plurality of normative demands'⁴³, then the corporation is permitted to act, indeed required by law to act, in a manner that is profoundly *irresponsible*.

By design, the corporation is incapable of responding to a plurality of normative demands. Rather, the corporation responds to a *single* normative demand – advancing the financial interests of the corporation and, by extension, its shareholders. In practice, this response takes the form of an automatic, asymmetrical distributive function, whereby the corporation distributes the financial rewards of corporate conduct *internally*, and distributes the costs of corporate conduct *externally*. These externalised costs may include risks of harm, financial or otherwise, including grave risks. As it derives from the corporation's monist, self-referenced normative posture, the corporation's distributive function may be characterised as inherently irresponsible.

7. Corporate irresponsibility – implications for the regulatory task

If it is to be at all effective in realising intended distributions of benefits and risks associated with AI, regulation must be properly sensitised to the corporation's inherent and

⁴⁰ S. Bottomley, *The Responsible Shareholder*, Edward Elgar, Cheltenham 2021, pp. 175-180.

⁴¹ P. Ireland, *Limited Liability, Shareholder Rights and Corporate Irresponsibility*, cit., p. 845.

⁴² Ireland describes this mix of control rights with the 'no-obligation, no-responsibility, no-liability' nature of corporate shareholding as 'a recipe for irresponsibility'. Ivi, pp. 853-854.

⁴³ G. Williams, *Responsibility as a Virtue*, cit., p. 459.

irresponsible distributive function⁴⁴. Faithful performance of the regulatory task requires conceiving of the task as a sub-set of corporate regulation, where the primary goals are anticipating and guarding against corporate irresponsibility.

Reframing the regulatory task in this way carries important implications for performance of the task. One of these implications is that harm mitigation approaches are a poor choice for AI regulation. This claim may be made in circumstances where the nature of the regulatory task is as described immediately above, and also as it is described in Part 3. The basis for this claim is that harm mitigation approaches are especially vulnerable to obstruction by corporate irresponsibility, at least insofar as they seek to distribute risk in a manner that differs to the corporation's own intrinsic distributive function. Given the monist design of the corporation's intrinsic distributive function, such differences are likely.

As noted in Part 3 above, the regulatory task tends to be directed towards two overarching concerns: 1) promoting innovation in AI; 2) and mitigating against, or providing redress for, AI harms. Further, as noted above, harm mitigation style regulation tends to establish these two concerns as antagonistic to one another; the less mitigation there is, the more innovation there is, etc. As innovation equals reward and mitigation equals cost, this antagonistic framing is easily hijacked by the corporation's intrinsic, irresponsible distributive function.

Consider for example, the *EU AI Act* (the Act). A combination function/risk-based instrument, in part the Act operates by classifying AI systems into different risk categories with attendant risk mitigation responsibilities. The most onerous risk mitigation obligations attach to systems categorised as 'High Risk' (HR). Interestingly, depending on the system, providers and deployers may play a role in determining whether a system is HR or not⁴⁵, potentially exempting themselves from the mitigation obligations that would otherwise attach to the system⁴⁶. Where a system is classified as HR, providers play a role in shaping their mitigation obligations. For example, under article 9, providers of HR systems must create and maintain a 'risk management system', which involves, amongst other things, identifying risks that the provider must then attempt to mitigate. Deployers also have a role to play in shaping their obligations under the Act, including in identifying risks associated with their systems, and identifying 'serious incidents' relating to the use of their systems⁴⁷.

Even from this brief overview the flaws in this approach should be apparent. Risk mitigation approaches routinely require delegation to regulated entities, including corporations⁴⁸. This involvement creates an opening for the corporation's intrinsic distributive function to displace the intended risk distributions of the regulatory instrument. Where providers or deployers are corporations, their intrinsic distributive function will incentivise under-classification of HR systems, understating of possible risks, and under-reporting of serious incidents, amongst other things. Doing so will alleviate costs and promote corporate financial gain⁴⁹. The likelihood that corporations would seek to

⁴⁴ Cfr J. Black & R. Baldwin, *Really Responsive Risk-Based Regulation*, cit., pp. 188-197.

⁴⁵ See for example *EU AI Act* article 6(3).

⁴⁶ For an insightful critique of this 'loophole' see S. Wachter, *Limitations and Loopholes in the EU AI Act and AI Liability Directives: What this Means for the European Union, the United States and Beyond*, «Yale Journal of Law & Technology», 26, 3, 2024, pp. 684-686.

⁴⁷ *Eu AI Act* article 26(5).

⁴⁸ J. Black & R. Baldwin, *Really Responsive Risk-Based Regulation*, cit., p. 201.

⁴⁹ Cfr Ivi, p. 202.

minimise or even evade mitigation obligations is enhanced by the obvious difficulties associated with detecting such behaviour, not least as ‘conformity assessments’ under the Act are undertaken by providers themselves, a mechanism aptly described by Wachter as ‘a major legal loophole’⁵⁰. Ultimately, the Act is insensitive to corporate irresponsibility and highly vulnerable to being obstructed by it as a consequence.

8. *A possible solution*

While it is beyond the scope of the present discussion to examine in detail how the problem of the corporation’s inherently irresponsible distributive function might be overcome, it is worth mentioning in closing one possible solution. Arguably, the key to sensitising AI regulation lies in harmonising the different aims of such regulation. Regulators should view the twin concerns of promoting innovation in AI and mitigating potential AI related harms as being mutually reinforcing, rather than antagonistic, considerations. That is, regulators should approach the concept of innovation as *entailing* safety. Entangling the considerations in this way may help to confound the corporation’s irresponsible distributive function. Disturbing the corporation’s ability to treat innovation and risk mitigation as zero-sum propositions will necessarily impede the corporation’s ability to externalise risks of harm, particularly where corporations are the developers of AI technology. Conceiving of innovation as entailing safety would train regulatory attention away from risk mitigation strategies, and toward approaches based on positive legal obligations to produce safe outcomes as a matter of fact. It will involve distributing the risks of AI harms back towards the developers and deployers of AI, building these costs back into the costs of development in particular. While they are hardly perfect models, product safety regulation and aviation regulation offer promising inspiration⁵¹, as these areas tend to involve positive legal obligations to produce safe outcomes, rather than positive legal obligations aimed at merely reducing risk⁵².

⁵⁰ S. Wachter, *Limitations and Loopholes*, cit., p. 692.

⁵¹ Cf B. Judge et al, *When Code Isn’t Law: Rethinking Regulation for Artificial Intelligence*, «Policy and Society», 44, 1, 2025, pp. 87-89. See also the *Australian Consumer Law*, and provisions such as section 54 which requires providers to ensure that goods and services are of an ‘acceptable quality’, including by being fit for purpose, acceptable in appearance, free from defects, safe and durable. Strictly speaking, section 54 operates as a guarantee, rather than a positive obligation, however the practical effect is quite similar, and hence it is interesting legislative artifact when viewed from a design inspiration perspective.

⁵² This would likely involve adopting a ‘zero risk approach’ (eliminating potential harms) or an ‘acceptable risk approach’ (reducing risks of harm to acceptable thresholds): see C. Coglianese, *The Law and Economics of Risk Regulation*, cit., p. 9.