Robot Rules: Regulating Artificial Intelligence
Jacob Turner
Palgrave Macmillan
2018
ISBN 9783319962344

Robot Rules is a timely publication addressing the current and emerging regulatory issues within the field of Artificial Intelligence (AI).

AI is a growing and dynamic field. While AI and its discussion are nothing new,[1] earlier discourse has chiefly been within the realm of science fiction.[2] Only recently has the discussion transitioned from the pages of science fiction to a modern reality (p 5). While works of fiction, such as Asimov's *I Robot*, have been an important starting point for discussion, in such a dynamic area, the loose, untested rules of fiction are inadequate. It is for this reason that a general overview of the challenges of AI and consideration of prospective regulatory solutions, as set out in *Robot Rules*, is particularly welcome.

Current discussion of the challenges of AI tends to focus on a specific legal problem at a specific point in time, perhaps best illustrated by contemporary discourse of AI in autonomous vehicles.[3] This has meant that in a dynamic industry, the general legal challenges going forward are only beginning to be discussed.[4] The fact that Turner's work does not confine itself to one strand of AI technology makes it particularly valuable, making it likely to retain relevance as AI technology develops.

The book is divided into eight chapters, concerning three main elements, namely the *rights* AI should have, *responsibility* for the AI's actions and the *ethics* applicable to the choices AI makes. Other publications have considered these aspects,[5] but none in as holistic a manner as *Robot Rules*.

Turner's main argument is that AI is unlike other new challenges which the law has faced. This is due to AI's dynamic nature, changing after creation, capable of independent agency. As a result, regulation must adapt to accommodate these changes. Turner does not seek to write those rules, nor is he calling for entirely new regulations. He does, however, provide a blueprint for consideration.

Chapter one starts by outlining the novelty of AI and its definitional difficulties.

1 The term AI has its origins in a summer workshop at Dartmouth College in 1956: Dartmouth, 'Artificial Intelligence (AI) Coined at Dartmouth' <250.dartmouth.edu/highlights/artificial-intelligence-ai-coined-dartmouth> (accessed 15 March 2020).

2 See the influential Three Laws of Robotics in Isaac Asimov, *I, Robot* (Doubleday, 1950) 40.

3 See for instance *Automated Vehicles: A joint preliminary consultation paper* (Law Com No 240, 2018); *Automated Vehicles: Consultation Paper 2 on Passenger Services and Public Transport: A joint consultation paper* (Law Com No 245, 2019).

4 *House of Lords Select Committee on Artificial Intelligence, AI in the UK: ready, willing and able?* HL 100 (2018). See also the ongoing White Rose research project led by P Morgan, *AI Law and Ethics: The Challenge of AI Wrongdoing* <https://whiterose.ac.uk/collaborationfunds/ai-law-and-ethics-the-challenge-of-ai-wrongdoing> (accessed 14 March 2020).

5 On rights, see D J Gunkel, *Robot Rights* (MIT Press, 2018). On responsibility, see J F Weaver, *Robots Are People Too: How Siri, Google Car, and Artificial Intelligence Will Force Us to Change Our Laws* (Praeger, 2013), 19–30 and R van der Hoven van Genderen, 'Legal Personhood in the Age of Artificially Intelligent Robots' in W Barfield and U Pagallo (eds), *Research Handbook on the Law of Artificial Intelligence* (Elgar, 2018). On ethics, see P Lin, K Abney and G A Bekey, *Robot Ethics: The Ethical and Social Implications of Robotics* (MIT Press, 2011).

While critical of the over and under-inclusiveness of current definitions, Turner does not seek to comprehensively redefine AI. Rather, he articulates *why* a specific and workable definition is essential, enabling prospective guidance and regulation. Turner's definition centres on AI's choice-making ability, leading to Turner to define AI as 'the Ability of a Non-natural Entity to Make Choices by an Evaluative Process' (p16).

Turner acknowledges that his definition is non-exhaustive, focusing on AI that is salient from a legal perspective. While this is justifiable, given the challenges of creating a workable definition for both present and future purposes, this omission is noteworthy for negligence lawyers, particularly medical negligence lawyers, as diagnostic AI may fall outside Turner's definition. This is because his definition excludes classical AI, the technology behind some medical diagnostics AI.[6]

Chapter two considers why AI presents unique challenges for regulation, demanding reconsideration of fundamental legal principles. The underlying argument is that legal systems are underpinned by agency and causation. Given that AI can make independent moral choices and through machine learning can develop independently, these factors challenge the current notions of legal agency and causation. In essence, the dynamic nature of AI and its capacity to develop without human input, generates unprecedented legal challenges.

Chapter three considers responsibility for AI, both in terms of liability for harm and positive outputs such as authorship in creative works, discussion of the former proving particularly insightful for negligence lawyers.

Turner considers a range of existing legal mechanisms which may be used to regulate AI. A common theme throughout this is the disruptiveness of AI to existing private law regulatory mechanisms. While the range of contractual and tortious regimes mentioned (negligence, strict liability, product liability, vicarious liability, no-fault compensation schemes, contract and third-party insurance) have positive aspects, the problems lie in their inability to accommodate AI fully.

The difficulty in shoehorning AI into existing private law mechanisms leads Turner to consider whether conferring AI with rights (chapter four) or granting legal personality to AI (chapter five) could be a solution. Thus, the shortcomings of current regulatory frameworks should be read in conjunction with Turner's discussion of legal personality for AI.

The European Commission has echoed these concerns,[7] particularly in terms of questions of liability.[8] Unlike Turner, the Commission have dismissed legal personality for AI as a solution of issues of liability, instead concluding that harm can be attributed to 'existing persons or bodies',[9] under current tortious and contractual rules.[10] Nevertheless, the Commission does not rule out legal personality for AI, acknowledging that

6 Two examples include R S Khan, A A Zardar and Z Bhatti, 'Artificial Intelligence based Smart Doctor using Decision Tree Algorithm' (2017) Vol 11 No 2 *ITB Journal of Information and Communication Technology*, and B Heinrichs and S B Eickhoff, 'Your Evidence? Machine Learning Algorithms for Medical Diagnosis and Prediction' (2019) *Hum Brain Mapp* 1.

7 European Commission, *Liability for Artificial Intelligence and other Emerging Digital Technologies Report from the Expert Group on Liability and New Technologies – New Technologies Formation* (2019).

8 European Commission, *Report from the Commission to the European Parliament, the Council and the European Economic and Social Committee Report on the Safety and Liability Implications of Artificial Intelligence, the Internet of Things and Robotics*, Brussels, 19.2.2020 Com (2020) 64 Final.

9 *Supra* n 7, 37.

10 *Ibid* 16.

it may be pursued if it helps to resolve emerging legal challenges.[11] Given the aforementioned regulatory challenges posed by AI, it is surprising that the Commission does not consider legal personality as an option, perhaps due to opposition to this proposal at an earlier stage.[12]

Although discussion of legal personality is not a novel suggestion,[13] as the discussion moves from the theoretical to the practical, consideration of legal personality as a solution becomes increasingly valuable.

In considering whether this is possible, Turner draws analogies with corporate law, highlighting the need to separate legal personality from the ability to make independent decisions, as with companies.

In considering whether this is desirable, Turner dismisses arguments from sceptics of legal personality for AI. According to Turner, the least tenable objections to granting AI legal personality are based on the 'mistaken conflation' of personality with humanity (p 189). To grant legal personality to AI is not to afford it all the same rights which humans enjoy.

Instead, it is the legal rights which would be granted, such as the ability to hold property, which are more complex and controversial. One such criticism of legal personality for AI comes from Bryson, Diamantis, and Grant[14] to which Turner refers frequently.

The argument is that AI is itself unaccountable and granting AI rights without obligations would make matters worse. To this, Turner asks, why robots cannot be given legal obligations? Indeed, giving AI the ability to hold property would allow it to settle debts and pay compensation in its own right. However, Turner does not elaborate further on *how* exactly AI could hold assets, and what would happen where there is no meaningful recourse against the AI. Given the reliance on Bryson et al in this section, and the fact this aspect is central to Bryson et al's criticism of granting AI legal personality,[15] it is surprising that Turner does not respond in detail. Furthermore, there is a risk that Turner's approach erroneously treats legal accountability and legal personhood as synonymous. To have the former does not require the latter, and *vice versa*.[16]

Chapter six looks generally at how we should '*design, implement and enforce*' bespoke rules made for AI (p 207). The legal vacuum resulting from governments failing to act has meant that private companies have made the first regulatory moves. Turner warns against reliance on self-regulation as it lacks legitimacy and impartiality. The voluntary nature of such rules mean they lack consistency and clout.

Given these problems, Turner argues that AI regulation should consist of a global code consisting of general rules applicable across industries. Indeed, as many of the novel challenges are the same across all sectors of AI (p 218), consistency and predictability are enhanced by such a solution. Drawing on several case studies, relating to internet domain names and space law, Turner seeks to demonstrate successful examples of international cooperation. While AI is

11 *Ibid* 38.

12 European Parliament Resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)), Article 59(f); Open Letter to the European Commission Artificial Intelligence and Robotics http://www.robotics-openletter.eu/ (accessed 19 May 2020).

13 L B Solum, 'Legal Personhood for Artificial Intelligences' (1992) 70 *NCL Rev* 1231.

14 J J Bryson, M E Diamantis, and T D Grant, 'Of, for, and by the People: the Legal Lacuna of Synthetic Persons' (2017) 25 *Artificial Intelligence Law* 273.

15 *Ibid* 288.

16 U Pagallo, 'Apples, Oranges, Robots: Four Misunderstandings in Today's Debate on the Legal Status of AI Systems' (2018) 376 *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 7.

distinguishable from these self-contained examples, these case studies demonstrate how international regulation could work, with incremental change, agreement on general principles and the involvement of international regulatory bodies.

Chapters seven and eight consider what ethical standards there should be for AI technology and how these can be applied. This is something on which the EU has consulted,[17] so Turner adds another voice to the growing discourse. Chapter seven focuses on controlling those who create the AI technology, considering how regulation and regulatory bodies for AI could operate in practice. Turner's underlying argument is that public confidence in the rules made is fundamental for such rules to be successfully implemented and adopted. The underlying and recurring principles are determining liability for harm caused by AI, safety in design of AI, transparency and requiring AI to operate in line with existing values. Turner suggests that licensing users to operate AI may be one of the ways in which harm from AI can be prevented. Chapter eight considers what rules can form 'minimum building blocks' (p 320), acting as open-ended and indicative suggestions for future regulations.

In sum, Turner's work is a valuable contribution to an under-discussed field. The criticisms noted are minor, and the book is a very welcome addition given the dearth of practical, accessible discussion on the topic. Turner's approach to facing the challenges raised by AI is prospective, rather than retrospective; general rather than specific. This means the book is likely to remain relevant despite future developments in AI technology.

What makes Turner's publication so important is that it concisely discusses the novel legal problems of AI in a manner accessible to those without expertise in either law or computer science. Turner draws analogies using a vast range of cultural and literary examples to illustrate the complex arguments at a level easily understood by the non-expert. This accessibility does mean some aspects of the work are not discussed in as much detail as desired, but this is justifiable in an overview text. Furthermore, the vast range of sources referenced means the reader is directed to other materials to further their learning.

Robot Rules is an engaging and accessible work which will be of interest to anyone concerned with the regulatory questions posed by AI.

Katie McCay
University of Bristol

17 European Commission, High-Level Expert Group on Artificial Intelligence, *Ethics Guidelines for Trustworthy AI* <https://ec.europa.eu/futurium/en/ethics-guidelines-trustworthy-ai/stakeholder-consultation-guidelines-first-draft> (accessed 12 March 2020).

Vicarious Liability: Critique and Reform
Anthony Gray
Hart Publishing
2018
ISBN 9781509920235

The law of vicarious liability has been 'on the move'[1] for some time now. There can be little doubt that the 'move' has all been in one direction, and consisted exclusively of an expansion of the circumstances in which vicarious liability arises. Only recently has there been any indication that the expansion may finally be slowing down.[2] Whatever the courts' motivation (or indeed the merits of their motivation) for this expansion, the result has been that it is now almost universally agreed that the current law of vicarious liability has no satisfactory rationale. As much is explicitly acknowledged by the courts.[3] Such a clearly unsatisfactory state of affairs has been the subject of much academic criticism – criticising the cases has simply become too easy.

In *Vicarious Liability: Critique and Reform*, Anthony Gray, Professor of Law at the University of Southern Queensland, seeks to bring some order to this chaos. In particular, he examines and critiques the various theories that are often said to justify the imposition of vicarious liability. Most of Gray's criticisms are directed at the currently fashionable 'enterprise risk theory of vicarious liability', which, he argues, is 'weak' (p x). He then outlines his 'preferred model of conceptualising the liability of an employer for actions of their employees', being the 'agency theory of vicarious liability', which holds that 'an employer should [only] be liable for actions of their employee agent which were for the purposes of the employer and for their benefit' (p x). According to Gray, his agency theory of vicarious liability 'will produce what I consider to be better outcomes than the current state of affairs' (p xi).

The book consists of three main parts, and 11 chapters in total. Part one, which consists of four chapters, provides a detailed overview of the law of vicarious liability – after all, if we are to critique the law, we must know what it is first (p 3). Chapter one starts by outlining the historical development of the doctrine from Roman times up to the early twentieth century. The next three chapters outline the development of the doctrine in the UK, Australia, and North America (Canada and the USA). The common theme throughout the chapters is the inability of the examined jurisdictions to articulate a satisfactory rationale that underpins the law, but that the idea of 'enterprise risk' certainly appears to be the favoured explanation.[4]

1 *Various Claimants v Catholic Child Welfare Society* [2012] UKSC 56, [2013] 2 AC 1 [19] (Lord Phillips).

2 *Barclays Bank plc v Various Claimants* [2020] UKSC 13, [2020] 2 WLR 960; *WM Morrison Supermarkets plc v Various Claimants* [2020] UKSC 12, [2020] 2 WLR 941.

3 See, for example: Lord Clyde in *Lister v Hesley Hall Ltd* [2002] 1 AC 215 [35] ('I am not persuaded that there is any reason of principle or policy which can be of substantial guidance in the resolution of the problem of applying the rule in any particular case'); and Lord Pearce in *Imperial Chemical Industries Ltd v Shatwell* [1965] AC 656, 685 ('The doctrine of vicarious liability has not grown from any very clear, logical or legal principle but from social convenience and rough justice').

4 Gray is clear that this is less so in the High Court of Australia, which has instead preferred explanations based on 'policy' (see, for example, Gummow and Hayne JJ in *New South Wales v Lepore* (2003) 212 CLR 511) – though, without articulating what that 'policy' is – or ignored the problem altogether (see, for example, the commentary to this effect in J Goudkamp and J Plunkett, *Vicarious liability in Australia: on the move?* (2017) 17 OUCLJ 162, 168, regarding *Prince Alfred College Inc v ADC* [2016] HCA 37).

Part two of the book, which also consists of four chapters, is the meat of the book. It outlines and critiques the various 'theoretical rationales that seek to justify or explain the doctrine of vicarious liability' (p 101). Chapter five focuses on explaining the 'enterprise risk' theory of vicarious liability, which Gray states is based on the economic rationale that 'a particular enterprise should have costs allocated to it that fairly reflect the cost of doing business' (p 111). In other words, if enterprises want the *benefits* associated with having employees, they must also accept (and pay for) the *costs* associated with having employees. The chapter then (ostensibly) moves onto 'scholarly support for enterprise risk theory', but much of the section in fact discusses *criticisms* of enterprise risk as a justification for vicarious liability, notwithstanding that this is supposed to be the focus of the *next* chapter.

In any event, Chapter six does indeed focus principally on critiquing enterprise risk theory. Gray's criticisms include: enterprise risk was not the historic justification for the imposition of vicarious liability; enterprise risk does not explain the existing contours of the law of vicarious liability, including why liability only attaches to employees but not independent contractors, nor why the imposition of vicarious liability is dependent on proof the employee committed a wrong; it is not clear why an enterprise risk rationale should be confined to the law of vicarious liability, when it plays no role in, and is arguably even inconsistent with, the rationale of other areas of tort law (most notably the common law's approach to product liability); it makes 'brave' (ie unrealistic) assumptions about ability of enterprises to spread and insure against certain losses; the ultimately economic justification does not explain how judges do decide cases nor how they should decide cases; and it fails to deter employers from causing harm. He concludes that courts should 'reconsider their use of the theory' (p 148).

Chapter seven briefly discusses and dismisses as unconvincing a number of other (ie non-enterprise risk) rationales that have also occasionally been suggested as justifying the law of vicarious liability, including the difficulty claimants may face in obtaining evidence of negligence, that employers are more likely to have 'deep pockets', and that it is simply more 'fair, just, and reasonable'.

Chapter eight outlines Gray's primary thesis, being that an 'agency theory' to vicarious liability is the way forward. According to his agency theory, vicarious liability will only attach where:

> 'the worker concerned had actual or apparent authority to do the thing for which the engager is now argued to be (vicariously) liable. Here, questions of whether the employee was acting for the employer's purposes and benefit will be crucial … Only in such cases can morally blameworthy behaviour be attributed to the employer so as to justify liability.'[5]

Gray says that such an approach *can* justify the imposition of vicarious liability because, 'a person who is in control of the actions of the person doing the wrong is considered sufficiently "morally blameworthy" to be responsible for the wrong' (p 163).

Gray acknowledges that under his approach victims of sex abuse would necessarily be left without a remedy, as 'in no sensible way could it be argued that an employee's … sexual abuse of a child is for the benefit of the employer and in furtherance of their interests' (p 172). Another implication of Gray's stricter approach is that there would be no liability

5 P 159.

in cases where an employee commits an act that was 'specifically forbidden by their employer' (p 174); after all, he says:

> 'the law generally allows parties to bargain about liability … [so] why is an employer not free to bargain with their employee, by in effect saying to them that these are the parameters of our contractual working arrangement … If you go outside those limits, you may well be personally liable for what you have done, but I will not be.'[6]

Whilst this would presumably allow employers to avoid ever being vicarious liable for employees' torts by simply inserting a term into their employment contracts along the lines of 'The employee will not be negligent', Gray believes his stricter agency model of vicarious liability is desirable because:

> 'The law of tort should not make liable someone who is not morally blameworthy. An employer who employs someone to do something and tells them not to do it in a particular way, should not be liable if the employee in fact does it in that way, causing injury to another.'[7]

Gray then argues that under his agency model of vicarious liability, employers would not be liable for the tortious conduct of independent contractors. In particular, he says, 'I am not liable because they are not my employee. I do not control their actions. As such, I have committed no morally blameworthy act for which I should legally be held responsible' (p 185). Accordingly, whilst an employer would be liable for the actions of a personal chauffer, they would not be liable for the actions of a taxi or uber driver (p 185).

Part three of the book, which consists of three chapters, is devoted to 'Miscellaneous Issues'. These issues are, the relationships to which vicarious liability might attach (ie employee vs independent contractor), non-delegable duties, and the circumstances in which vicarious liability should permit the award of punitive damages. The issues are not examined through the lens of Gray's agency theory of vicarious liability, but at a more generalised level, meaning that the discussion should still be of interest even to those who (like me) were not persuaded that his agency approach is the way to go.

Vicarious Liability: Critique and Reform is a well-researched book which clearly shows that the law of vicarious liability is a complete mess. Aside from being a useful up-to-date overview of the seemingly ever-changing law, it is particularly valuable for its examination of vicarious liability from a more foundational level. As Gray notes, it is plainly unsatisfactory that there is currently no clear rationale for the doctrine's existence at all. Gray is therefore surely right that, given the increasing influence and popularity of the enterprise risk theory of vicarious liability, an 'in-depth analysis … is warranted' (p x). Gray succeeds in providing that in-depth analysis which ultimately raises some serious questions about the ability of enterprise risk to justify the current law, and shape it moving forward.

Of course, as is inevitable in a book of this nature, there are also a number of aspects that are unlikely to appeal to everyone. In my view, two aspects of the book are particularly deserving of comment.

First, whilst Gray certainly highlights a number of difficulties the enterprise risk theory faces when trying to explain the law

6 P 174.
7 P 176.

of vicarious liability, these seemed (at times) somewhat overstated. So, for example, the fact that the *historical* justification for vicarious liability was not enterprise risk theory is surely no reason it cannot (or should not) justify the *current* law. Similarly, whilst it might pose problems for those searching for a unified theory of tort law, the fact that enterprise risk theory is not currently used to justify (or shape) other areas of tort law does not mean that it cannot provide a satisfactory justification for the current state of the law of vicarious liability. As for the potential difficulty employers may face in obtaining insurance or spreading loss, this has no bearing whatsoever on the ability of enterprise risk theory to explain the law; it is a practical issue at best. And the fact that enterprise risk fails to deter enterprises from causing injuries (at least over and above what the law of negligence already deters) is not a criticism of the theory at all; it merely establishes that deterrence is not a convincing argument in *favour* of it (just as countless other bad arguments are not either). There is, of course, still the problem that enterprise risk theory fails to explain the current law, and particularly why vicarious liability only attaches to employees but not independent contractors, and why the imposition of vicarious liability is dependent on proof the employee committed a wrong. This is indeed a serious problem for the enterprise risk theory of vicarious liability.[8] Yet, with the requirement that the wrongdoer is an employee of the defendant no longer a necessity,[9] it is certainly not the problem it used to be, such that enterprise risk theory would seem to come much closer to providing a satisfactory explanation of vicarious liability than Gray suggests.

Second, I was not convinced that Gray's agency theory of vicarious liability was tenable. The immediate problem is that if the law is to make a principal vicariously (ie strictly) liable for the wrongdoing of his agent, this *also* needs to be somehow justified, even in cases where the agent/employee is acting for the 'benefit or purposes' of his principal/employer. Pointing out that this is what the law of agency already does doesn't help – it merely shifts the problem elsewhere. And if the rationalisation is something along the lines of, 'principals get the benefits of their agent's actions and so must also incur the costs',[10] then that sounds a lot like enterprise risk theory, which Gray has already dismissed as 'weak'. Gray anticipates this objection and so clarifies that an agent/employee will only be acting for the 'benefit or purposes of' his principal/employer where the latter has 'authorised, expressly or implicitly, the employee to commit a wrong' (p 178); that is, 'authorised an action that is recognised by the law as a tort' (p 179). The basis for Gray's narrow approach to agency is that, 'only employers who are *themselves* morally blameworthy (emphasis added)' should be liable for employees' actions; 'the law should not impose liability on a person or organisation that is not at fault' (p 178). Gray concedes that some may object that if this is what he means by 'agent', then it is not true vicarious liability at all, but an example of *direct* liability. Gray responds that he would 'not object to this change in terms' (p 178). Strictly speaking, then, Gray's response does indeed succeed in justifying his agency model of vicarious

8 A point I have made previously: J Plunkett, 'Taking Stock of Vicarious Liability' (2016) 132 *LQR* 556.

9 See, for example, *CCWS*, *supra* n 1, *Cox v Ministry of Justice* [2016] UKSC 10, [2016] AC 660, and *Armes v Nottinghamshire County Council* [2017] UKSC 60, [2018] AC 355. Of course, the Supreme Court has recently clarified that, whilst the wrongdoer does not need to be an 'employee' (ie the relationship only needs to be 'akin to employment'), vicarious liability still does not extend to 'independent contractors': *Barclays Bank*, *supra* n 2.

10 Whilst Gray does not specifically endorse this explanation, a similar rationale seems to come across in many of the cases he cites.

liability independently of enterprise risk; the employer is 'vicariously' liable because they are *directly* liable. Such an answer, of course, raises a number of even more fundamental difficulties for Gray's agency theory of vicarious liability. First, it throws the baby out with the bathwater; if tortious liability requires moral blameworthiness (ie fault), then Gray is presumably against *all* strict-liability torts (eg the rule in *Rylands v Fletcher*) and doctrines that don't require personal fault (eg the objective standard of care in negligence). This apparent implication is significant and surely requires further explanation. Second, it effectively renders Gray's agency theory entirely unnecessary; a principal is 'vicariously liable' for the actions of his agent (under Gray's model) because he is *directly* liable – the fact that the wrong was committed via an agent is seemingly irrelevant. Third, whilst new theories of law will never be a perfect 'fit' with the existing law, Gray's 'agency theory of vicarious liability', which ultimately seems to be little more than axing the doctrine of vicarious liability entirely, doesn't fit with the current law at all, and in far more fundamental ways than the enterprise risk theory of vicarious liability. Finally, it is ultimately intellectually unsatisfying; it would be like proposing to resolve the long-running debate about whether the law of tort is based in rights/corrective justice or distributive justice by abolishing the law entirely – that way it would be neither!

My criticism of Gray's book should not, however, be overstated. Gray has set himself an almost impossible task in seeking to single-handedly bring some order to this famously difficult area of law; that (at least in my view) he has not done so is hardly surprising. It should also be noted that Gray's agency theory only forms a small part of *Vicarious Liability: Critique and Reform*. As noted above, his summary of the existing law is detailed and helpful, as is his identification of where the law has gone wrong. Gray's analysis of the enterprise risk theory of vicarious liability is also timely and much needed (even if a bit overstated), given the extent of the influence it appears to have on modern courts. Ultimately, *Vicarious Liability: Critique and Reform* makes a welcome contribution to the law of vicarious liability.

Dr James Plunkett
Barrister
Victorian Bar