RIGHTS, SOCIAL JUSTICE AND RESPONSIBILITY IN THE LAW OF TORT

EMMANUEL VOYIAKIS*

I INTRODUCTION

Harry carelessly runs Sally down in the street, causing her serious physical harm. Consider two propositions about their situation:

1. Harry is responsible in tort towards Sally, because he has violated her right to physical safety, thereby committing a wrong against her.
2. The fact that Sally may be the beneficiary and Harry the victim of an unjust distribution of resources should make no difference to Harry's tortious responsibility towards her.

I will assume that both propositions are true. Many theorists take that truth to entail that what is typically called 'distributive justice' is not a proper normative standard by which tort law ought to be assessed. I want to suggest that this idea is mistaken for two related reasons. The first reason is that the idea misunderstands its target. Theories of 'distributive justice' are not theories about distribution. They are theories about the basic structure of society. The second reason is that social-structural concerns play an important role in determining the general conditions of moral responsibility, and are therefore important elements in explaining when the violation of a right makes an agent responsible to the bearer of that right. If this is correct, theorists who exclude considerations of social justice from the normative standards that apply to tort law not only misunderstand the nature of social justice and its demands, they also have a poor account of responsibility for violations of rights.

II TORT LAW, RIGHTS AND DISTRIBUTIONS

It is a widely held view that the aim of tort law, or the normative standard to which tort law should be held, is to provide redress against infringements of rights – and the correlative duties that agents have towards one another – and not

* Lecturer, Department of Law, London School of Economics and Political Science.

to effect a just distribution of resources across the community.[1] Advocates of this 'rights thesis' do not deny that just distribution is a legitimate social aim. They only claim that achieving that aim is not the job of tort law, but of other parts of the law, such as our tax and welfare systems.

One might object that even if tort law might not have a distributive *aim*, it has a distributive *impact*.[2] Tort law rules determine which classes of agents will bear the cost of accidents and other harm-causing events that occur when those classes of agents interact. For example, they determine when businesses will bear the cost of accidents caused by their products and suffered by consumers, and when employers will bear the cost of the torts of employees. Furthermore, it is clear that changes to the relevant tort law rules will typically entail changes in the distribution of rights and duties between those classes of agents; to that extent, such changes will have an impact on the overall position of agents relative to one another. Other things being equal, a regime of negligence liability for defective products will leave consumers to shoulder more costs than a regime of strict liability would, an exclusively negligence-based regime of employers' liability will leave claimants (employees or third parties) in workplace accidents to shoulder more costs than a regime of vicarious liability would, and so on. If tort law rules have a significant impact on how much agents falling within a certain class have relative to agents falling in another class, then tort rules are unavoidably in the business of resource distribution, and therefore must be measured against the demands of whatever scheme of principles tells how we ought to distribute resources in a just society.

This objection does not work. The mere fact that a legal rule has a distributive impact does not of itself mean that we may legitimately tinker with that rule to achieve a certain distributive pattern.[3] We cannot say to Sally that she ought to pay for her own medical bills because Harry has a lot less than a just distribution would give him. Equally, Harry may be entitled to more, but he is not entitled to get it by having Sally pay for the harm he caused to her. The natural explanation of why such a way of shifting of costs and benefits around would be unjustifiable is that Sally has a *moral right* to her physical safety, Harry has a correlative duty not to infringe that right and that, by running into her, Harry has committed a wrong against her.[4] The upshot of that explanation is that

1 For some leading statements of that view, see Jules L Coleman, *Risks and Wrongs* (Cambridge University Press, 1992) ch 18; Ernest J Weinrib, *The Idea of Private Law* (Harvard University Press, 1995) ch 1; John C P Goldberg and Benjamin C Zipursky, 'Torts as Wrongs' (2010) 88 *Texas Law Review* 917; John Gardner, 'What Is Tort Law For? – Part 1: The Place of Corrective Justice' (2011) 30 *Law & Philosophy* 1; Scott Hershovitz, 'Two Models of Tort (and Takings)' (2006) 92 *Virginia Law Review* 1147; Robert Stevens, *Torts and Rights* (Oxford University Press, 2007) ch 2; Allan Beever, *Rediscovering the Law of Negligence* (Hart Publishing, 2007) ch 2.

2 Peter Cane, 'Distributive Justice and Tort Law' [2001] *New Zealand Law Review* 401.

3 William Lucy, *Philosophy of Private Law* (Oxford University Press, 2007) 340.

4 Cf Coleman, above n 1, 305: 'If sustaining or protecting a less than fully just distribution of wealth or resources can sometimes be a matter of justice, it cannot be a matter of distributive justice. Then what sort of justice is it that permits, if it does not explicitly, endorse distributive injustice? The answer is, corrective justice'.

no distribution can be justified if it fails to protect certain moral rights that agents possess, and to remedy the wrongs constituted by the infringement of those rights. While it is true that tort law rules have a distributive impact, they develop that impact in the course of protecting moral rights; most clearly, our right to our physical health and safety, and our right to our property.[5] It follows that we could not hope to justify a distribution that purported to deny protection to those rights in order to fit a broader distributive aim.

Note that the rights thesis need not deny that some parts of tort law might, in fact, be justified by reference to considerations of just distribution. In particular, advocates of that thesis could accept that the choice between different ways of protecting rights and/or remedying their infringement might turn on distributive considerations.[6] For example, the imposition of strict rather than negligence-based product liability might be justified on the ground that producers are better able to pass on the cost of liability insurance to the market, thus helping distribute the cost of the harm more fairly across the community. This flexibility is possible because the notion that a claimant has a moral right (and the defendant a correlative moral duty) may require that the state provide redress to the claimant against the defendant, but it does not necessarily specify exactly what form that redress should take, nor does it exclude non-right considerations from bearing on that further question. The core claims that the rights thesis needs to maintain in order to exclude just distribution from the normative standards applicable to tort law are: first, that the reason tort law protects the interests it does is that those interests correspond to certain rights that agents have; and second, that the content of those rights is independent of and normatively prior to considerations of just distribution.

Consider now a second objection to the rights thesis. That thesis claims that tort law is not in the business of distribution, because it protects moral rights that no scheme of distribution could afford to throw into the same pot with non-right considerations. The apparent implication here is that those who think that tort law rules are subject to the demands of just distribution fail to acknowledge that agents' rights place normative limits on possible distributions. This implication is demonstrably mistaken. John Rawls' theory of justice provides an obvious and familiar example. That theory is built on the idea that all agents ought to enjoy certain rights, and that those rights are lexically prior to principles that distribute basic goods. Regardless of whether one agrees with Rawls about what rights those are, and how their lexical priority should be understood, the general setup of his theory clearly allows one to accommodate both of the core claims of the rights thesis: it is possible to say that tort law is part of a scheme of just distribution, and still maintain both that the reason tort law protects certain interests of agents is that those interests correspond to moral rights, and that those

5 For present purposes, I leave aside the question whether we have a moral right to the protection of certain of our economic interests.

6 Coleman goes further, accepting that considerations of corrective justice may sometimes be trumped by other applicable considerations, including those of distributive justice and cost efficiency: see Coleman, above n 1, 392–5. For a view that denies this possibility, see Weinrib, above n 1, ch 7.

rights are normatively prior to other distributive considerations. If that is correct, then accounts of tort law that find its justification in a theory of just distribution would have a distinct explanatory advantage over the rights thesis: they would be able to account for the place of *both* right and non-right considerations in the justification of tort law rules.[7]

This objection also fails, because it is no objection to the rights thesis at all. Its basic upshot is that the rights thesis can be consistent with a theory of distributive justice. However, advocates of the rights thesis can take this idea in their stride, because it challenges neither of their core claims: that tort law rules are justified insofar as they protect moral rights, and that those rights are normatively prior to considerations of just distribution. After all, advocates of the rights thesis trade precisely on the idea that the business of distribution cannot get started before we have secured certain moral rights that agents have.[8]

It seems to me that the two challenges to the rights thesis fail because they pit the idea of just distribution *against* the idea of agents' moral rights. I want to suggest that the relationship between them is far less adversarial than either the rights thesis or the objections to it acknowledge. My suggestion will depend on two shifts of focus. The first shift moves from the idea that theories of justice like Rawls' are theories about distribution, to the idea that they are theories about the design of the basic *social structure*. The second shift entails a move from the idea that infringements of a claimant's moral rights constitute wrongs, to the idea that infringements of a claimant's rights make an agent morally *responsible* to that claimant. In this way, I hope to draw attention to a little-noted desideratum for any theory that traces the normative foundations of tort law in the notions of moral rights and wrongs: that any such theory should be consistent with our considered convictions about the limits of moral responsibility, and about the impact of social structures on that responsibility.

III FIRST SHIFT: FROM DISTRIBUTION TO THE SOCIAL STRUCTURE

Talk of justice and its relation to tort law sets most lawyers and theorists thinking about the distinction between 'distributive' and 'corrective' justice. That distinction may be important, and the stakes of deciding which (if any) of those two kinds of justice private law should pursue may be significant, but the debate is somewhat hampered by the fact that the name 'distributive' gives a bad characterisation of the kind of justice it is attached to. Calling a kind of justice 'distributive' encourages the impression that what is special about 'distributive justice' is that it distributes burdens and benefits and, by extension, that there are

7 For this view, see Kevin A Kordana and David H Tabachnick, 'On Belling the Cat: Rawls and Tort Law as Corrective Justice' (2006) 92 *Virginia Law Review* 1279.

8 On that point, see Peter Benson, 'The Basis of Corrective Justice and Its Relation to Distributive Justice' (1992) 77 *Iowa Law Review* 515; Jules Coleman, 'The Practice of Corrective Justice' in David Owen (ed) *Philosophical Foundations of Tort Law* (Oxford University Press, 1995) 53, 66–7.

kinds of justice (for example, 'corrective justice') that do not involve distribution. That suggestion is misleading. Principles of corrective justice also distribute burdens and benefits when they say that victims of wrongs are entitled to the benefit of compensation on the part of wrongdoers, who must bear the corresponding (typically economic) burden. The real difference between these various kinds of justice lies not in the fact that one distributes and the other does not, but in the *reasons* that underlie the distributions that each produces. By that token, the names that we give to each kind of justice ought to tell us something about the character of the justification given by each for the distribution that it recommends. The name 'corrective justice' fits this bill well; it tells us that a certain distribution is justified by the fact that one agent committed a wrong against another and a principle that requires our institutions to correct it. The name 'distributive justice' does not; it gives us no clue as to the character of reasons that ought to guide the distributions that this kind of justice requires and how those reasons might differ from reasons of corrective justice.

Calling that other kind of justice *social* justice is a much better option. If any theory qualifies as one of what private law theorists call 'distributive justice', John Rawls' certainly does. However, the Rawlsian (and the Marxist) conception of justice takes its subject matter to be not distribution *simpliciter*, but distribution produced by the *basic structure of society*, and holds that a social structure is just when the resource distributions that it produces could be justified to the agents that are subject to them.[9] More simply, a Rawlsian conception of social justice does not care about how much people have (or don't have) by way of resources; it cares about how much social structures cause people to have (or not to have) by way of resources. Whether or not we accept the particular principles that Rawls thought apply to the basic structure of society, taking that structure as the subject matter of social justice has a distinct consequence for the relationship between social justice and tort law. It shows that the real question is not whether tort law should be sensitive to reasons relating to resource distribution, but whether it should be sensitive to reasons that apply to the basic structure of society; or, to put it more plainly, whether we have reason to use tort law *to effect a certain design on that social structure*.[10] To say that tort law should do social justice is to answer this question in the affirmative.

---

9 John Rawls, *A Theory of Justice* (Harvard University Press, rev ed, 1999) 36–40. Rawls' particular way of appealing to this contractualist idea was famously based on the mechanism of the 'original position'.

10 Cf Dennis Klimchuk, 'On the Autonomy of Corrective Justice' (2003) 23 *Oxford Journal of Legal Studies* 49, for an example of a definition of distributive justice that makes no mention of the role of the basic structure:

> distributive justice is that domain of justice concerned with the distribution of benefits and burdens among members of a given group, who enjoy the relevant benefits and shoulder the relevant burdens either simply owing to their membership in that group, or in accordance with some measure of entitlement which applies to them in virtue of their membership: at 50.

Klimchuk's definition would include in the domain of distributive justice the distribution of a birthday cake amongst family members. I have no stakes in showing that cake distribution could not be regarded as involving some 'localised' sort of distributive justice; my concern is to highlight the difference between 'localised distributive' and 'social' justice.

For present purposes, I will take the basic structure as consisting of the complex web of facts and relations that affect the interaction of agents across the community and across time, in respect of subject matters of interaction that are central to agents' overall life projects and prospects, such as employment, consumption, transport, and so on.[11] This understanding is close to the broader conception of the basic structure that some of Rawls' own writings advocate, in the sense that it is not limited to the institutional structure that guarantees basic liberties and oversees the large-scale distribution of resources across the community (mostly through tax-and-transfer systems).[12] That broader conception relies on three intuitions. First, that the background conditions against which individuals interact matter for the purpose of assigning responsibilities to the parties, and of deciding the stance that our institutions should take towards such interaction. Second, that those background conditions include the basic configuration of a community's economy and social organisation, inasmuch as that configuration shapes some important roles that an agent may occupy in the context of typical and important interactional contexts (for example, entrepreneur or employee, business or consumer, driver or pedestrian), and the life prospects that are associated with being able to interact with others from within that role.[13] Third, that agents able to modify the basic structure to their or someone else's advantage (institutions certainly fall in this category) ought to justify such modifications to all affected agents within that structure. Accordingly, I will treat reasons that 'apply to the social structure' as being reasons that recommend a certain design of the basic structure by virtue of the fact that such a design could be justified to all agents subject to it. While none of those intuitions is uncontroversial, I will not say anything to shake the convictions of those who

11 Here I am broadly following A J Julius, who defines the basic structure as 'a distribution of the population over types and an assignment of situations to types that are together reproduced by the distribution of actions they induce', where *situations* are the list of factors relevant in the explanation of how an agent behaves in conditions of intentional interaction and *types* are groups of people in similar situations: A J Julius, 'Basic Structure and the Value of Equality' (2003) 31 *Philosophy & Public Affairs* 321, 330. See also T M Scanlon, *What We Owe to Each Other* (Harvard University Press, 1998) 244: 'The "basic structure" of society is its legal, political, and economic framework, the function of which is to define the rights and liberties of citizens and to determine a range of social positions to which different powers and economic rewards are attached'.

12 Kordana and Tabachnick discuss the 'narrow' and the 'broad' conception of the basic structure in Rawls' work and demonstrate, first, that Rawls' writings prevaricate between this narrow and a broader view of the basic structure, and second, that the broader view is compelling for those who accept Rawls's two principles of justice: see Kordana and Tabachnick, above n 7. As my argument will not rely on Rawls' principles of justice or take extended issue with his conception of the basic structure, their arguments will be only occasionally available to me.

13 Arthur Ripstein, 'Public Order and Private Justice: Kant and Rawls' (2006) 92 *Virginia Law Review* 1391, 1395 ff offers a conception of the basic structure premised on the Kantian idea that 'people are required to forbear from interfering with one another' and that 'provided they do so, the only grounds of co-operation are voluntary': at 1407–8. The notion of 'forbearing from interfering with one another' seems to me to suggest that if everyone has their fair share and everyone stays out of each other's way, except when voluntarily transacting with them, any resulting distribution will be just. That would not be correct; the very idea of a basic social structure is precisely the higher-order explanation of how iterated interactions between members of a community across time and subject matters may impact on other agents' interactional options, even when those agents have *never* interacted directly with one another.

would take issue with them and, accordingly, I will make no effort to sell the rest of my argument to that (largely, but not exclusively, libertarian) constituency.

Although much contemporary thinking about social structures and their impact on individual transactions recommends that a social structure is just when it meets certain standards of equality,[14] I will not assume the truth of that claim. I will assume only that such justifications should be broadly contractualist in character. In particular, following Scanlon's version of contractualism,[15] I will suppose that social justice obtains when no agent could reasonably reject the principles that assign each agent their particular place within the basic structure and hold each responsible for the ways in which he or she interacts with others against the background of that structure.[16]

Against that background, we can frame the suggestion that tort law should aim to do social, rather than distributive, justice as follows. The structure of our societies makes it necessary for classes of people to interact with each other from within certain roles in order to secure their wellbeing. Some of those interactions take place in settings in which one class of agents imposes risks on the other on a non-reciprocal basis (for example, businesses impose risks on consumers, employers impose risks on employees, and so on). Other interactions involve the imposition of more reciprocal or near-reciprocal risks (for example, driving, or the common use of one's land).[17] In both cases, rules that determine who bears the cost should those risks materialise are themselves part of that social structure, because they constitute an aspect of the background conditions of interaction between those classes of agents across the community and across time. Therefore, those rules fall to be assessed in the light of reasons that apply to the social structure itself, namely reasons of social justice.

The shift from distribution to the social structure as the subject matter of justice does not pose a direct challenge to the rights thesis. Advocates of that thesis could still maintain, with much plausibility, that protecting certain rights is normatively prior to implementing a certain social structure, and therefore limits what may count as a legitimate structural design. In the next two Parts I will nevertheless try to show that focusing on the social structure rather than on distribution presents the rights thesis with an indirect but no less potent challenge.

---

14 See Matthew Seligman, 'Luck, Leverage and Equality: A Bargaining Problem for Luck Egalitarians' (2007) 35 *Philosophy & Public Affairs* 266.

15 Scanlon, above n 11, ch 5.

16 I will also not stake a claim as to whether reasons that apply to the basic structure of society so conceived are reasons that involve duties for institutions only, or whether they involve practical duties for individuals too (the duty to support just institutions aside). For the latter view see G A Cohen, 'Where the Action Is: On the Site of Distributive Justice' (1997) 26 *Philosophy & Public Affairs* 3; Liam B Murphy, 'Institutions and the Demands of Justice' (1998) *27 Philosophy & Public Affairs* 251, 258.

17 George P Fletcher, 'Fairness and Utility in Tort Theory' (1972) 85 *Harvard Law Review* 537; Gregory C Keating, 'Fairness and Two Fundamental Questions in the Tort Law of Accidents' (Olin Working Paper No 99-21, University of Southern California Law School, January 2000) <http://papers.ssrn.com/sol3/papers.cfm?abstract_id=200778>.

IV SECOND SHIFT: FROM MORAL RIGHTS AND WRONGS TO MORAL RESPONSIBILITY

Suppose that we ask not what moral rights Sally has or what moral wrong Harry has committed against her, but whether Harry should be morally responsible to Sally for causing her harm. It is tempting to think that the move from moral rights to moral responsibility tells us less than we already knew. If anything, the fact that Sally had a right to her physical safety and that Harry had a correlative duty not to infringe that right *explains* why Harry should be responsible towards Sally for the harm he caused her. Phrasing the question in terms of responsibility seems to tell us less precisely because it does not reveal the ground of that responsibility. The omission is not trivial: an agent may be morally responsible to others for some choice or outcome even when no interpersonal rights and duties are involved. I am responsible for making you sad because I let my begonias, which I knew you loved so much, wither away and die, even though you had no right to require me to keep taking care of them and I had no duty to you to do so. However, my moral responsibility differs from Harry's in an obvious practical sense. Sally's moral right entitles her to demand that Harry behave in a certain way. Your interest in me maintaining my beautiful begonias comes with no such entitlement; it is merely a *pro tanto* but defeasible consideration that I ought to take into account in deciding what to do or how to behave towards you. Telling the story in terms of moral rights and duties seems better, because it allows us to keep the difference between the two scenarios in the foreground of our attention.

It would nevertheless be a mistake to think that stories about moral responsibility simply sit on top of stories about moral rights and duties. If the violation of a moral right is a particular way in which an agent may become responsible to the bearer of that right, then the general conditions of moral responsibility will also be conditions of responsibility for the violation of a moral right. For example, suppose that our preferred theory of moral responsibility says that an agent may only be held responsible for outcomes that are the consequence of his or her choices. That would entail that the violation of a right would make an agent responsible to the bearer of the right only as long as that violation was a consequence of the agent's choice. Or suppose that we agree with Joseph Raz that an agent is morally responsible for the outcome of some conduct only when that conduct is the result of the functioning (successful or failed) of the agent's powers of rational agency.[18] It would follow that the violation of a right would make an agent responsible only if the conduct that constituted the violation was the result of the agent's successful or failed rational functioning.

The point is not that whenever we should refrain from holding Harry responsible for harming Sally, Sally's moral right to her physical safety is not

18 Joseph Raz, 'Responsibility and the Negligence Standard' (2010) 30 *Oxford Journal of Legal Studies* 1, 15.

violated. It is, rather, that any institution that holds agents accountable for causing harm to others needs to be justified by reference to a *theory of moral responsibility* rather than just a theory of rights. Whatever else it might be, tort law is clearly such an institution, and therefore the fact that Harry might not be morally responsible for the harm he caused to Sally may well make a difference as to which liabilities tort law may legitimately assign to him. Under the rights thesis itself, tort law is not there to tell us what moral rights agents have, but to tell us how the state ought to *respond* to violations of those rights.[19] For that thesis to succeed, it must be the case that violations of rights entail moral responsibility for the violating agent. But this claim is not straightforwardly true. Perhaps an agent who violates a moral right should not always be held responsible for that violation (for example, when the agent was coerced into violating the right). Or perhaps moral responsibility should sometimes lie with an agent other than the agent who violated the right (for example, when employers have to bear responsibility for workplace accidents caused by their employees). In any event, the rights thesis must rely on some theory of responsibility for violations of rights and, in turn, any such theory should be consistent with what we take to be the general conditions of moral responsibility.

Given that there are several ways of deploying the idea of responsibility in general and of moral responsibility in particular, I want to distinguish the sense in which my discussion uses that idea.[20] For a start, I will assume that ascriptions of *legal* responsibility ought to be capable of being justified towards agents, and that judgments of *moral* responsibility will be an important part of any such justification. Furthermore, the judgments of moral responsibility I have in mind are *substantive* judgments about the practical burdens that we may legitimately require an agent to take on when we try to formulate general principles for the regulation of our conduct towards each other. Substantive judgments of this sort, as I will understand them, are distinct from judgments about the identity of the agent or agents that a certain course of conduct may be *attributed* to for the purposes of moral assessment ('it was X who engaged in this conduct'), and judgments about the *causal* impact of that conduct on another agent's interests ('it was X who caused Y's harm').[21] At the same time, I will not assume that judgments of substantive moral responsibility entail judgments about moral *blameworthiness*; in other words, I will not take the fact that an agent is morally responsible for some outcome as entailing that other agents are justified in

19 The point applies even if one is thinking about legal rather than moral rights; the point of tort law cannot be to tell us what legal rights agents have, but to tell us what forms of redress violations of those rights entitle their bearers to: see Goldberg and Zipursky, above n 1, 947–52.

20 For a general discussion of the senses in which ascriptions of responsibility may be relevant in the law, see H L A Hart, *Punishment and Responsibility: Essays in the Philosophy of Law* (Oxford University Press, 2nd ed, 2008) 211–31; Peter Cane, *Responsibility in Law and Morality* (Hart Publishing, 2002) ch 2.

21 We might say that the first sort are judgments of *attributive* responsibility, and the second sort are judgments of *causal* responsibility; Hart treats both under the rubric of causal responsibility: Hart, above n 20, 214–15.

adopting an attitude of blame towards that agent.[22] Finally, I will not stake a view as to whether an agent's substantive moral responsibilities derive from the practical roles that the agent occupies vis-a-vis others, or whether they may emerge ad hoc.[23] In short, my discussion will assume that we have answered questions of attribution and causation,[24] and will focus on the question of substantive responsibility. For the sake of simplicity, I will also not dwell on the downstream problem of how to apportion substantive responsibility when the relevant principles dictate that more than one agent ought to have a share in it.

In the remainder of this Part, I want to outline an account of substantive moral responsibility that I find plausible, drawing on Scanlon's idea of responsibility as a function of the 'value of choice' for an agent.[25] I believe that the substantive moral claims that follow from that idea can be derived from more than one overall account of morality, but for present purposes, I will only appeal to its intuitive force. I will go on to claim in Part V that those who find the 'value of choice' account plausible have reason to reject a core claim of the rights thesis, namely that one can determine the content of an agent's rights without regard to social-structural considerations.

The idea that there is a close connection between responsibility and choice is a central feature of our intuitions about responsibility. One way to describe that connection is to say that we have good reason to want our responsibilities to depend on our choices. Wanting this is not the same as wanting our responsibilities to depend on our having complete control over what happens. In normal circumstances, an agent does not have control over how the dice will come up, but this is not a ground on which he or she can disclaim responsibility for the result once he or she has chosen to throw them. Rather, we see luck and risk as built into the nature of certain options available to the agent, and treat the degree of luck-dependence or the riskiness of an option as a substantive factor to be taken into account in determining the responsibilities of an agent who chooses

22 Hart seemed to take a contrary view. He wrote that 'to say that a person is morally responsible for something he has done or for some harmful outcome of his own or others' conduct, is to say that he is morally blameworthy, or morally obliged to make amends for the harm': ibid 225. I follow Scanlon in thinking that moral blameworthiness and moral responsibility to make amends towards (or compensate) another agent are distinct, in that the latter is a necessary but not a sufficient condition of the former: T M Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame* (Harvard University Press, 2008) ch 4.

23 For a critical discussion of Hart's concept of role-responsibility in that regard, see Cane, above n 20, 34–6.

24 For present purposes, I elide some difficulties about the relation between these two senses of responsibility. One such difficulty is that attributing conduct to agents requires some account of how conduct should be described, and it seems clear that substantive moral considerations should play a part in that. For example, the fact that Harry's employer may be one of the persons that principles of attributive responsibility pick out as being a fit subject of moral assessment seems to be explained by the fact that principles of substantive responsibility sometimes require that we hold employers responsible for the wrongs of their employees.

25 Scanlon, above n 11, ch 6.

it.[26] We get closer to a good description of the relation between responsibility and choice by saying that we find value in having our responsibilities depend on how we respond to situations that present us with different options, and the existence of reasonably favourable conditions for deciding which option to choose. On our original scenario, we can plausibly say that Harry is responsible for harming Sally because (or insofar as) he had the option of driving with more care, and faced no impediment or other condition that made it exceedingly difficult for him to take that option. If we look at the matter *ex ante*, we could say that a principle that ascribes responsibility to an agent is justified when that agent could have avoided responsibility by choosing appropriately.

Pointing to the value of having our responsibilities depend on our choices is, of course, just a starting point. Any theory that places choice at the heart of moral responsibility will have to give an account of what we might call the *circumstances of choice*. An intuitive way of approaching the issue is to insist that choice entails responsibility when that choice is 'free' or 'voluntary', but it seems to me that the intuition is no more than a placeholder for more particular moral concerns about the *range of options* available to the agent, and of the *conditions* under which the agent is called upon to choose amongst those options.[27] That is to say, an agent is likely to have a reasonable objection to being held responsible for some choice, either on the ground that the agent's range of options was excessively or improperly constrained, or on the ground that the conditions under which the agent had to choose amongst those options were such that it was too difficult for the agent to choose appropriately.[28] Had Sally thrown herself in front of Harry's car before he could react, Harry would have had an objection of the first sort against being held responsible for Sally's harm, while had Harry crashed into Sally because that was the only way to avoid an oncoming truck that was driving in the wrong lane, he would have had an objection of the second sort. Debates about whether Harry's choice in either example is 'free' or 'voluntary' are really conclusions of arguments about the

26 For a well-known deployment of the idea that the presence of luck does not preclude ascriptions of responsibility to an agent, see Tony Honoré, 'Responsibility and Luck' (1988) 104 *Law Quarterly Review* 530; John Gardner, 'Obligations and Outcomes in the Law of Torts' in Peter Cane and John Gardner (eds), *Relating to Responsibility: Essays for Tony Honoré on His Eightieth Birthday* (Hart Publishing, 2001) 111, 126–8.

27 Compare Raz's discussion of the significance of valuable options in the context of individual autonomy: Joseph Raz, *The Morality of Freedom* (Oxford University Press, 1988) 374–5.

28 Some philosophers have argued that Scanlon's account does not state sufficient conditions for an agent's responsibility; the fact that the availability of choice might lead agents *generally* to choose appropriately does not explain why agents who are differently disposed should be responsible for having chosen differently: see Andrew Williams, 'Liberty, Liability and Contractualism' in Nils Holtug and Kasper Lippert-Rasmussen (eds), *Egalitarianism: New Essays on the Nature and Value of Equality* (Oxford University Press, 2007) 241. Others have argued that Scanlon's value of choice account is too restrictive in this connection; having a choice may be valuable for an agent even when the agent is placed in a situation where he or she is very likely to choose badly: see Alex Voorhoeve, 'Scanlon on Substantive Responsibility' (2008) 16 *Journal of Political Philosophy* 184. As neither objection challenges the two aspects of the value of choice theory that I develop below, I will not take them up in the present context.

seriousness and moral significance of these sorts of constraints on his choice, rather than independent moral standards.

Naturally, we can expect that people will disagree as to when such objections to responsibility may legitimately be put forward. All choices we make are circumscribed in one way or another; it is hardly ever the case that we find ourselves with all the alternative options we could ask for, and ideal conditions to ponder the choice among them. Determining how constrained a choice must be for that choice not to make an agent responsible is a substantive moral question, and the way in which the answer to it is structured will depend in large part on one's overall moral theory. Contractualists like Scanlon will approach it by asking whether the agent could reasonably reject a putative principle that took the range of options that the agent had available, and the conditions under which he or she had to choose, as sufficient grounds for holding that agent responsible for the outcome of his or her choice. Utilitarians will ask whether a principle of that content would be optimific, and so on. However, it seems to me that we can draw attention to two particular aspects of the significance of choice for the purposes of moral responsibility without having to settle on one grand moral theory. The first aspect is that principles that treat moral responsibility as a function of an agent's choice involve a division of labour between the risks and responsibilities to be undertaken by the choosing agent, and those to be undertaken by other agents (and, sometimes, the community as a whole). The second and related aspect is that, when proposing principles of responsibility and designing institutions that will hold agents to account in accordance with those principles, the challenge is not simply to ensure that agents' responsibilities depend not only on their choices, but to ensure that we are doing enough to put those agents in a reasonably favourable position for choosing what to do.

Scanlon elaborates on the first aspect as follows:

> Because most people take themselves to be more actively concerned with the promotion of their own safety and well-being than others are, they want outcomes to be dependent on their choices even when this has only 'avoidance value.' Given this concern, 'giving people the choice' under favorable conditions makes it extremely unlikely that they will suffer easily avoidable harms. We do not want the trouble and expense of supervising others' choices more closely, and do not want them to be supervising us. Therefore, we take the view that giving people the opportunity of avoiding a danger, under favorable conditions, often constitutes 'doing enough' for them: the rest is their responsibility.[29]

At the same time, wanting our responsibilities to depend on our choices, on the ground that we take ourselves to be in the best position to choose what will promote our safety and wellbeing, implies that the corresponding division of responsibility between agents should be sensitive to the fact that, in certain contexts and under certain circumstances, we are likely to choose poorly, even in the presence of safe alternatives. When that is the case, we may reasonably want

29 T M Scanlon, Jr, 'The Significance of Choice' in Sterling M McMurrin (ed), *The Tanner Lectures on Human Values* (Cambridge University Press, 2011) vol 8, 149, 197. The ideas presented by Scanlon in his 1986 Tanner Lectures here and at nn 30–31 below came to form the basis of Scanlon, above n 11, ch 6.

some degree of outside supervision or external systems of support for harm avoidance, especially if these could be provided without excessive cost to others:

> Even if we know that actions avoiding a certain unwanted outcome will be available to us in a given situation, we also know that our processes of choice are imperfect. We often choose the worse, sometimes even in the knowledge that it is the worse. Therefore, even from the point of view of an agent looking at his own actions over time, situations of choice have to be evaluated not only for what they make 'available' but for what they make it likely that one will choose. It is not unreasonable to want to have some protection against the consequences of one's own mistakes.[30]

The second aspect of the significance of choice for responsibility is related to this last point. Given that we are not always in a good position to assess alternative options and to act upon them, we have reason to opt for principles that ascribe responsibility to an agent, and to design institutions that hold agents accountable in accordance with those principles, when we have done enough to place that agent in a reasonably favourable position for choosing. The idea that the presence of choice is itself a sufficient justification for holding an agent responsible

> exaggerates the importance of the fact of choice relative to that of the conditions under which the choice was made. [It] suggests that these conditions are important only insofar as they bear on the voluntariness of the choice. This is a mistake. The fact that a choice was voluntary does not always establish that we 'did enough' for an agent by placing him or her in the position from which the choice was made. Nor does the fact that an agent did not voluntarily choose an outcome, or choose to take a certain risk, establish that what resulted was not his fault. Giving him the *opportunity* to choose may have constituted 'doing enough' to protect him.[31]

The point is that the *exercise* of choice on the part of an agent has derivative significance in the context of assessing the agent's responsibility. When the agent was placed in sufficiently favourable circumstances for making a choice, its exercise will be only the last piece needed to complete the jigsaw puzzle of responsibility, and will owe its significance to the fact that the rest of the pieces are already in place.[32] When, on the other hand, the circumstances of choice are wrong, the fact that the exercise of that choice was voluntary cannot by itself justify holding the agent responsible, insofar as we have not done enough to place the agent in the position that he or she could reasonably want to be in when choosing.

We can test these intuitions in further variations of the scenario involving Harry and Sally. Suppose that Harry had suffered a heart attack whilst driving and lost control of the car, crashing into Sally. Or suppose that Harry was a newly qualified driver, and that, through poor design, the tests that he had to pass in order to obtain his driving license had failed to examine him on certain aspects of driving skill. Suppose, finally, that Harry would have avoided running into

30 Scanlon, above n 29, 195. This also explains why Raz's rational functioning principle does not state a sufficient condition for moral responsibility.

31 Ibid 196 (emphasis in original).

32 Cf Scanlon, above n 11, 260–1.

Sally if his car had been equipped with a state-of-the-art collision avoidance system.

In the first variation, Harry's crashing into Sally was not the result of a choice on his part. However, that fact may not necessarily provide Harry with an objection to a principle that held him responsible for Sally's injuries. Under the value of choice account, what matters is whether Harry was placed in a good enough position to assess the risks involved and to act upon that assessment. Harry's possible knowledge that he could be suffering from a heart condition would be an important consideration here, precisely because such knowledge (or its absence) changes the position of Harry relative to that of every other agent in the scenario. If Harry was already being treated for a heart condition, he would probably have been in a better position to estimate the danger that he might suffer heart attack whilst driving, and would therefore have no objection to a principle that made him responsible for choosing to undertake that risk.[33] If Harry had no reason to suspect that his heart might fail him, then the value of choice account instructs us to ask whether we could reasonably require of Harry that he assume the risk of suffering an unknown disease whilst driving. In turn, that will require us to look into the frequency of such incidents, their typical severity and the suddenness of their onset, any knowledge of predisposition towards the disease that Harry might have had, and so on.[34]

The second scenario requires us to ask whether the failures of the testing system had the effect of putting Harry in an unreasonably poor position from which to choose how to drive. The value of choice account suggests that answering this question will require us to estimate the importance of the aspects of driving skill that the driving test omitted to test, and their relevance to the accident that Harry caused, as well as any limitations to the range of dangerous situations that one might reasonably expect driving tests to cover. If the test failed to examine Harry on his ability to maintain steady course in normal traffic, then Harry could reasonably complain that the official certification of his driving ability made him a source of risk for others, while lulling him into a false sense of security about his skills. Part of the point of having a public licensing system is precisely to ensure that all drivers meet certain basic standards of skill, and agents have reason to want to be able to rely on the results of that test in assessing whether they meet that standard. At the same time, the range of situations to which a driving test can expose a candidate will necessarily be limited, as it may be difficult to create realistic simulations of certain real-life dangers, so Harry could not reasonably require that the test have examined every facet of his driving ability.

In the third scenario, determining the effect of the availability of a state-of-the-art collision avoidance system on Harry's responsibility will require us to look into the affordability of that system, its reliability, and the ease of operating

33 Cf *Roberts v Ramsbottom* [1980] 1 WLR 823, in which the defendant was found to have been in breach because he had knowledge of his medical condition and had suffered similar attacks in the past.

34 Cf *Mansfield v Weetabix Ltd* [1998] 1 WLR 1263, in which the defendant was held not to be in breach, given that their driver had no reason to anticipate the condition.

it. One notable point in this scenario is that the enquiry into the value for Harry of being able to choose in these circumstances is not exhausted by a consideration of Harry's own driving skills. Even if no competent driver would have been able to avoid the collision by relying on their driving skills, Harry may still be responsible for the accident if his choice not to install the system, and to undertake the corresponding risk that he may be involved in a road accident that driving skill would not have been able to deter, was made under favourable conditions. For example, if the system was very cheap to purchase and install, reliable, and easy to operate, the value of choice account would suggest that Harry would not be able to point to the fact that he exercised all due driving skill to avoid responsibility for Sally's harm. The fact that he was placed in a good enough position to take measures to avoid such collisions would be sufficient to make him responsible for the resulting harm.

I have argued that the rights thesis could only hope to work as a normative account of tort law if it supplies us with a theory of responsibility for violations of moral rights, and that, in turn, such a theory should be consistent with the general conditions of moral responsibility. I hope I have also shown that seeing responsibility as a function of the value of an agent's choice has the advantage of pointing our attention to features of situations involving harms caused by one agent to another that do not register adequately under a rights-based account. These are the range of options available to the agent who caused some harm, and the circumstances in which that agent had to choose between those options. The fact that questions about what sort of conduct agents are entitled to demand from one another make it harder to bring those features into proper focus is not a coincidence. Questions about rights are typically asked with a view to describing a state of affairs that ought to obtain in favour of the bearer of the right, and to giving a canonical and practically usable formulation of that agent's entitlements towards others. Questions of responsibility, on the other hand, are more concerned with an agent's *judgment* about the applicable reasons; they are typically asked when the world is messier and we are in search of the agent that ought to clean it up. Although stories about rights and duties and stories about responsibility should eventually dovetail, the difference between these two perspectives seems to fulfil a concrete practical purpose. We do not normally care to describe our rights beyond a certain level of specificity; for example, we are generally content to say that Sally has a right that others do not threaten her physical safety, and we do not qualify that right for every situation in which the agent who threatens Sally's physical safety ought not to be held responsible for doing so. The reason, it seems to me, is both that producing a full specification would be a practically impossible and therefore useless task,[35] and that we are taking it for granted that a reflective agent should not take a 'shorthand' canonical formulation of the right as an exhaustive statement of all the reasons

35 See Judith Jarvis Thomson, *The Realm of Rights* (Harvard University Press, 1990) 82 ff; Joel Feinberg, *Rights, Justice and the Bounds of Liberty: Essays in Social Philosophy* (Princeton University Press, 1980) 221 ff.

that might be involved in a given situation. However, when we get to the stage of ascribing responsibility and designing institutions that will impose specific burdens on agents on account of their conduct, the task of assessing all the applicable reasons confronts us directly, and considerations that formed the silent backdrop to the abstract formulation of the right come to the foreground of our attention.

V RESPONSIBILITY AND SOCIAL STRUCTURES

Let us now turn to whether our assessment of an agent's responsibility by reference to the circumstances of that agent's choice ought to take into account the state of the social structure. The abstract answer to that question is, I think, obvious. The background conditions against which agents interact – conditions that include the structures that make it necessary for certain classes of agents to interact from within certain roles, and the rules that determine who bears the cost in case the risks of harm associated with those interactions materialise – have a direct impact on the range of options available to agents in a given interactional setting, and the circumstances under which these agents get to choose among those options. When the conditions in question fail to place agents in the position they could reasonably want to be in when making the relevant choice, a principle that made responsibility contingent on those conditions could not be justified to those agents. This connection between social structures and responsibility should not be surprising. To aim for a just social structure is, in part, to aim for structures of social cooperation that place agents in conditions such that agents will be generally responsible for their choices in particular interactional contexts. The point is sometimes phrased in terms of a just social structure leaving agents 'free' to pursue their own ends in their interactions with others,[36] without having to worry about structural questions, but this clearly implies that responsibility about one's interactional choices in a just social structure is the other side of the same coin. Or, getting the social structure right is dealing with one source of potential objections that agents might have to being held responsible for their choices in certain interactional contexts.

It is still important to hash this abstract argument out in terms of concrete examples in which the social structure bears on an agent's responsibility by constraining his or her interactional options, both in order to give more definition to the relationship between responsibility and social justice, and to get a better handle on the operation of structural constraints on an agent's interactional options.

---

36 Cf John Rawls, *Political Liberalism* (Columbia University Press, 1995) 268–9:

> What we look for, in effect, is an institutional division of labor between the basic structure and the rules applying directly to individuals and associations and to be followed by them in particular interactions. If this division of labor can be established, individuals and associations are then left free to advance their ends … secure in the knowledge that elsewhere in the social system the necessary corrections to preserve background justice are being made.

Consider, for a start, an example involving voluntary rather than accidental interaction.[37] Our social structure requires agents occupying certain roles to enter into voluntary transactions with each other in order to secure their wellbeing, while placing one of those classes of agents in a significantly stronger bargaining position. Voluntary transactions between businesses and consumers, employers and employees, and insurers and insurance holders are typical and familiar instances of such structural imbalances in bargaining power. In these types of transactions, the options of the structurally weaker class of agents are constrained, sometimes in far-reaching ways, while those of the structurally stronger class are correspondingly expanded. Although the resulting imbalance might turn out to be perfectly justifiable to the weaker class as a general matter (for example, on the ground that it allows economies of scale and brings down costs for all consumers), it also allows structurally stronger parties to commit structurally weaker parties to clauses that have no value for the latter class, taking away bargaining options that this class might reasonably want to have available, either because those options are intrinsically important for it or because they are only worth trading for much higher increases in purchasing power. Against this background, the statutory regimes that bar the enforceability of some such clauses in consumer, employment or insurance agreements limit the responsibility of the structurally weaker class for having agreed voluntarily to those clauses.[38] Those regimes are justified precisely because they recognise that, say, a consumer's choice to agree to a wide range of clauses in a standard form contract with a business may not make the consumer responsible under those clauses, insofar as that choice has been made under objectionable social-structural constraints. And by changing the conditions of enforceability for certain types of voluntary transactions across subject matters and across time, those regimes become themselves part of the (now justifiably modified) social structure.

The shape of the social structure can have an impact on an agent's responsibility in the context of accidents too, either by constraining or expanding an agent's interactional options, or by requiring that agent to choose among available options under poor conditions. Liability for carefully produced but defective products is a familiar case in point. To determine whether producers or consumers should bear the risk that defective products might cause harm to consumers, we should ask whether such a principle of responsibility would allow either party to avoid those harms, or to minimise the cost of addressing them, by choosing appropriately. This will require us to take account of familiar factors, such as the seriousness of the threatened harms, the possibility of easy intermediate inspection by the consumer, and so on.[39] But the structure of our societies will also bear on this question in three distinct ways. First, the way our

37 I discuss such examples in Emmanuel Voyiakis, 'Contract Law and Reasons of Social Justice' (2012) forthcoming 25 *Canadian Journal of Law and Jurisprudence*.

38 In the United Kingdom, most notable among those regimes are the *Unfair Contract Terms Act 1977* (UK) c 50, and the *Unfair Terms in Consumer Contracts Regulations 1999* (UK) SI 1999/2083.

39 See, eg, *Grant v Australian Knitting Mills* [1936] AC 85.

economies are organised makes it necessary for all agents to be consumers in order to pursue their wellbeing. Second, the economies of scale on which modern industrial production is based make it practically impossible for a producer to check every product with utmost thoroughness, especially when the product in question has a short life cycle or is in massive circulation. Third, the same economies of scale ensure that the marginal increase in the price of a product when the cost of liability is built into production costs is likely to be reasonably small.[40]

Taken together, these considerations seem to me to offer a plausible explanation of why we may legitimately put responsibility for defective products on the producer rather than the consumer. Consider, in particular, the choice that a principle that imposed responsibility on consumers would be presenting them with. As giving up consumption and taking individual protective measures are not viable options, consumers will often be unable to avoid the risk of suffering serious harm from defective products. Furthermore, they may only minimise the cost of those harms to them by taking out health insurance, an option that will be less valuable to an agent the worse off that agent finds himself or herself in the social distribution of resources.[41] On the other hand, a principle that imposed responsibility on producers would not necessarily allow them to reduce the risks in question, but it would allow them to pass on the cost of insuring against such responsibility to the market by treating it as a production cost and adjusting the price of the product appropriately. As long as all competing products in the market are produced with equal care and therefore expose consumers to similar risks, the marginal increase in the price of the product would be similar for all producers, and would not, by itself, put any of them at competitive disadvantage. Given that the statutory regimes that determine who will bear the cost of the harms caused by defective products are a part of the background conditions against which producers and consumers interact, opting for a regime of producer responsibility entails opting for a social structure that uses producers as *middlepersons* for spreading the cost of such harms to the community. This last feature also seems to me a central part of explaining why imposing responsibility

40 *Escola v Coca-Cola Bottling Co of Fresno* 150 P 2d 436, 440–1 (Traynor J) (Cal, 1944):

> Even if there is no negligence, however, public policy demands that responsibility be fixed wherever it will most effectively reduce the hazards to life and health inherent in defective products that reach the market. It is evident that the manufacturer can anticipate some hazards and guard against the recurrence of others, as the public cannot. Those who suffer injury from defective products are unprepared to meet its consequences. The cost of an injury and the loss of time or health may be an overwhelming misfortune to the person injured, and a needless one, for the risk of injury can be insured by the manufacturer and distributed among the public as a cost of doing business.

41 Polinsky and Shavell have recently disputed the claim that liability for defective products lowers costs for consumers, partly on the ground that most consumers already have health insurance: see A Mitchell Polinsky and Steven Shavell, 'The Uneasy Case for Product Liability' (2010) 123 *Harvard Law Review* 1437. The point seems to me to ignore the fact that the holdings of an agent are a significant factor in determining the affordability of health insurance. On this point, see Didem M Bernard, Jessica S Banthin, and William E Encinosa, 'Wealth, Income and the Affordability of Health Insurance' (2009) 28 *Health Affairs* 887.

on producers does not constitute an objectionable constraint on their interactional options.

Regimes of vicarious liability offer another familiar illustration of the impact of the social structure on an agent's responsibility, though I will suggest that we have reason to doubt the justification of one particular aspect of those regimes in their current form. To return to our original scenario, suppose that Harry ran into Sally while in the course of doing business for his employer. In line with most jurisdictions, English law says that two agents are responsible for Sally's harm: Harry is responsible due to his negligent driving (negligence liability), and his employer is responsible too due to the relation between the activity Harry was carrying out and the employer's enterprise (vicarious liability). Now, it seems to me that the argument in favour of the imposition of responsibility on the employer proceeds along similar lines to the argument in favour of the imposition of responsibility on producers.[42] Employers, typically entrepreneurs, will usually be able to pass on the cost of the accidents caused by their enterprise to the market, through appropriate adjustments to the price of their product or service (therefore functioning as middlepersons for spreading the cost of enterprise liability to the community), whereas victims of such accidents may not be in a good position to protect themselves from the harms that the operation of enterprises may cause, or to obtain affordable insurance against the risk that such harms will materialise.

By contrast, consider the position of employees like Harry. To see whether the principle of negligence law that holds Harry responsible for the harm he caused to Sally in the course of his employment is justified, we should consider whether that principle allows Harry to avoid responsibility by choosing appropriately. On the face of it, the response seems to be obviously affirmative: Harry had the option of avoiding responsibility by driving carefully. However, I would suggest that stopping the analysis at this point would fail to register a number of further features of the circumstances of Harry's choice, features that relate directly to the structure of the employment context. For a start, our social structure makes it necessary for the majority of agents to take up dependent labour in order to secure their wellbeing. Most of such labour takes place in the context of complex enterprises, which very often create risks for employees and third parties, more so when the subject matter of the enterprise is an inherently dangerous activity. In turn, the complexity and inherent dangers of the enterprise have the effect of heightening the risks that may arise from careless choice on the part of an employee. Pressing the wrong button on a television remote control and pressing the wrong button on a factory floor involve vastly different risks and threaten harms of vastly different orders. Given that structural context, it seems reasonable for an employee to require some protection against the consequences of potential carelessness on his or her part, and to reject a principle

42 See Gregory C Keating, 'The Idea of Fairness in the Law of Enterprise Liability' (1997) 95 *Michigan Law Review* 1266; Douglas Brodie, 'Enterprise Liability: Justifying Vicarious Liability' (2007) 27 *Oxford Journal of Legal Studies* 493.

of responsibility that takes the mere fact of careless choice as a sufficient ground for holding him or her responsible for the resulting harm.

The fact that employees are practically *unlikely* to be pursued by the victim of a workplace accident does not, in my view, blunt the force of this objection, for two reasons. The first is practical: under current English law, employees remain exposed to the threat of liability, both against the victim, in the form of an action in negligence; and – more importantly – against employers and their insurers, in the form of an indemnity.[43] In that last regard, it should be noted that the 1959 Gentlemen's Agreement of the British Insurers Association – in which insurers undertake not to pursue indemnity claims against the employees of an insured employer except in cases of wilful conduct – is not a legally binding instrument, and at any rate does not cover all cases of workplace accidents.[44] The second reason is one of principle: if the justification for imposing negligence liability on employees for workplace accidents is weak, then the law ought to abolish it.[45] Furthermore, the recognition that negligence on the part of an employee should not make the employee personally responsible to the victim would also serve to clarify the nature of the employer's vicarious liability. In particular, it would demonstrate that vicarious liability is not some form of 'secondary' responsibility, or even a device for 'extending' responsibility beyond its natural bounds, but a direct or primary responsibility that follows from the position of enterprises in the structure of our societies.

If this suggestion stands up to scrutiny, it also throws into stark relief a serious problem for the rights thesis. There is no doubt that by driving carelessly and running into Sally, Harry violated Sally's right to her physical safety and committed a wrong against her. However, if responsibility for that wrong should not be shared between Harry and his employer, but should be borne exclusively by the latter, it would follow that the violation of a right – or the commission of a corresponding wrong – is neither a necessary nor a sufficient condition for holding an agent responsible for some harm. Or, to put the point conversely, the merits of the rights thesis as a normative account of tort law would emerge as conditional on the demands of a theory of responsibility, in which social structural considerations play an important part.

VI CONCLUSION

Tort law protects rights, and that protection cannot be legitimately compromised to achieve the aim of just distribution. This proposition may be

43 *Lister v Romford Ice and Cold Storage Co Ltd* [1957] AC 555. For an early criticism of employers' indemnities, see Ross Parsons, 'Individual Responsibility versus Enterprise Liability' (1956) 29 *Australian Law Journal* 714.

44 In *Morris v Ford Motor Co Ltd* [1973] QB 792, it was held that the 1959 Agreement did not cover harm caused by employees other than those of the insured employer.

45 Some Australian jurisdictions have: see, eg, *Employee's Liability (Indemnification of Employer) Act 1982* (NSW); *Law Reform (Miscellaneous Provisions) Amendment Act 1984* (NT).

true, but it encourages two important misunderstandings. First, the theories that tort lawyers refer to as theories of 'distributive justice' are not theories about distribution, but theories about the ideal design and justification of the *social structure* that shapes agents' interactional options across subject matters and across time. Second, the idea that tort law protects rights offers, at best, an incomplete normative account of tort law. Theories of tort law need to rely on some theory of *responsibility* for violations of rights; in turn, any such theory will need to be consistent with the general conditions of moral responsibility. The two misunderstandings are connected, because social-structural constraints on an agent's interactional options will often bear on the assessment of that agent's substantive moral responsibility, at least if we accept – as I have – that responsibility is a function of both the range of options available to an agent, and the conditions under which the agent is called to choose among those options.

I think that this argument also suggests that the dilemma whether tort law ought to protect rights (or provide civil recourse for their violation, or correct the wrongs such violations constitute), or whether it ought to maintain a just social structure, is greatly exaggerated. As an institution that ought to hold people accountable for the harms they cause, tort law is necessarily concerned with both. The obvious but perhaps underappreciated reason is that the rights and duties we have, and the structures within which we interact, both matter for determining the extent of our responsibilities towards each other.