Obligations and Outcomes in the Law of Torts (2001)

by John Gardner
Professor of Jurisprudence
University of Oxford
http://users.ox.ac.uk/~lawf0081

P Cane and J Gardner (eds), *Relating to Responsibility: Essays for Tony Honoré* (Oxford: Hart Publishing, 2001)

Obligations and Outcomes in the Law of Torts

JOHN GARDNER*

I

'Under strict liability,' according to Alan Brudner,

> one has no right to act in ways that happen to injure another. Since, however, all action carries the risk of such injury, strict liability means that I have a right that you be governed in all your actions by concern for my welfare, and you have the same right over me. No doubt there is a mutuality of care here; but it is the mutual care of extreme altruists who, because they claim no worth as independent selves, can neither give nor receive effective confirmation of worth and hence can require no valid right to care. By contrast a fault [i.e. negligence] requirement establishes a reciprocity of care between selves.[1]

Brudner casts his argument in terms of the potential plaintiff's rights, but (to avoid some distractions further down the line) I am going to reframe it in terms of the potential defendant's obligations. True enough, P's rights against D are not the same

* I am grateful to the participants in the Columbia Law School workshop on Tony Honoré's *Responsibility and Fault* for a fascinating discussion of my first attempt at the topic, from the ashes of which this rather different paper arose. Special thanks go to my commentator Niki Lacey and to Tony Honoré for giving me the benefit of his first reactions. Subsequent drafts were much improved by detailed comments from Peter Cane and Timothy Macklem.

[1] Alan Brudner, *The Unity of the Common Law* (Berkeley: University of California Press 1995), 190. For a similar line of thought, see Ernest Weinrib, *The Idea of Private Law* (Cambridge, Mass., Harvard University Press 1995), 182-3.

thing as D's obligations towards P. A right is not an obligation; rather, it is the ground of (one or more) obligations.[2] All the same, there is nothing that counts as the *violation* of a right other than a failure to perform (one or more of) the obligations that it grounds. Thus the conditions under which D violates P's rights – the conditions under which D *wrongs* P – are identical to the conditions under which D fails to perform (one or more of) his obligations towards P. It follows that, if our interest is in what it takes for D to wrong P, and in the liability to which such wrongs may give rise in law, the important differences between rights-talk and obligations-talk can largely be put on one side.

It seems to me, accordingly, that we can faithfully recast Brudner's argument in the following terms. When the law imposes strict liability on D for actions of D's that injured P, says Brudner, it asserts that D had (and failed to perform) an obligation to take extreme care not to injure P – such extreme care not to injure P, indeed, that D would have effaced himself entirely in performing the obligation. In demanding that D have acted so entirely self-effacingly for the sake of P, the law undermines its own position that P in turn owes the same obligation to take extreme care not to injure D. How could the existence of this reciprocal obligation be defended when, by the law's own reckoning, D clearly counts for so little?

This argument against the legal imposition of strict liability fails. I do not mean that it fails in its striking moral claims inspired by Kant and Hegel. Maybe Brudner is right that there comes a point at which one person takes such extreme care not to injure another that in the process she effaces herself entirely. Maybe Brudner is also right that there is something obnoxious about effacing oneself entirely in this way, or at least about being under an obligation to do it, or at the very least about being authoritatively held to be under an obligation to do it. And maybe that obnoxiousness comes, as Brudner suggests, of the fact

[2] Joseph Raz, *The Morality of Freedom* (Oxford: Clarendon Press 1996), ch 7.

that the self-effacing action is performed for reasons that cannot consistently be universalized to include the scenario in which D's and P's positions are reversed. Maybe Brudner is also right that, for this reason if no other, the law should not hold anyone to be under an obligation to take self-effacingly extreme care not to injure another. Let all of this be true. The problem remains that it still does not add up an argument against the legal imposition of strict liability unless it is also true, as Brudner asserts right at the start of the argument, that in imposing strict liability on D for actions of D's that injured P, the law holds D to have been under (and to have failed to perform) an obligation to take extreme care not injure P. And sadly this doctrinal premiss (as I will call it) is not true. In fact, it is an inversion of the truth.

The truth is that when the law imposes strict liability on D for actions of D's that injured P, it asserts that D had (and failed to perform) a straightforward obligation *not to injure P*. D would have performed this obligation not to injure P – and hence would have avoided the strictest of strict liabilities – if and only if she had not injured P. Now, not injuring P, and therein performing her obligation, is something that D might have done entirely fortuitously, without taking the slightest care, let alone extreme care, not to injure P. (Perhaps P just happens not to step off the pavement as D careers drunkenly and obliviously past in her car.) Conversely, injuring P, and therein failing to perform her obligation, is something that D might have done in spite of having taken the most extreme conceivable care not to do so, even to the point of total self-effacement that so worries Brudner. (Perhaps P has an unprecedented allergic reaction to all the cotton wool D spent the whole day laboriously wrapping him up in.) Under a regime of strict liability for D's injuring P there could be no liability in the first case where no care at all is taken, whereas there could be liability in the second where every conceivable care is taken. When strict liability is at stake, in other words, the law does not give two hoots *either way* about the care that D took. The most inhuman lack of care does not count

against her but by the same token the most superhuman investment of care does not count in her favour. It is this double-edged obliviousness to D's precautions that lawyers have in mind when they call liability by the evocative name 'strict'. So strict liability is not, as Brudner suggests, a variant of negligence-based liability in which the standard of obligatory care has been cranked up to an awesome level. It is a quite unrelated mode of liability in which the law does not care about care-taking, and therefore does not treat the bestowing of care – *any care at all* – as having been obligatory. Correspondingly, no measure of self-effacement at all is called for in performing the obligation.

You may object that in these remarks I am interpreting Brudner's doctrinal premiss uncharitably by interpreting it *too* doctrinally. I am attributing to him the clearly mistaken view that D's legal obligation, non-performance of which gives rise to D's strict liability, is an obligation to take extreme, self-effacing care not to injure P. But what Brudner is actually suggesting (you may say in his defence) is that D's declared legal obligation – admittedly a straightforward obligation not to injure P, for the performance of which taking care is admittedly neither necessary nor sufficient – effectively puts D under another (legally undisclosed) obligation as well, namely a morally obnoxious obligation to take extreme care not to injure P. To be exact, isn't Brudner's real point that taking the most extreme conceivable care, to the point of total self-effacement, is the *only available means* that D has to perform his straightforward legal obligation not to injure P? And that being so, doesn't the very logic of practical rationality dictate that D also has a second, derivative obligation to take that most extreme conceivable care? To will the end, after all, is to will the means; and to will that the end (D's not injuring P) be obligatory is surely, by the same token, to will the derivative obligatoriness of the only available means (D's taking ever-such-extreme care not to injure P). It is the obnoxiousness of this latter derivative obligation that, on this more charitable reading of Brudner, casts shame on the former

(parent) obligation, and hence casts shame on the strict liability that is based on the former obligation's non-performance.

I doubt whether this is really a more charitable reading. The proposed argument is open to challenge at virtually every step. To avoid tangential complications, however, let's concede one premiss of the argument straight away, namely the one ascribed to the very logic of practical rationality. According to this premiss (if I may express it as a general formula) an action γ is *derivatively obligatory* for a given agent D whenever another action ϕ is obligatory for D and γing constitutes D's *only available means* of ϕing. That premiss is relatively easy to grant, of course, because everything turns on what an *only available means* is. So what is it? The three words of the expression clearly signify three different conditions. To make the premiss a credible rendition of the logic of practical rationality, I propose the following as a minimal unpacking of the three conditions. For γing to be D's *only available means* of ϕing it has to be the case that

(a) D's γing would contribute to D's ϕing (it is a *means*);

(b) D has the capacity and the opportunity to γ in the circumstances (it is an *available* means); and

(c) no other action available to D would contribute to D's ϕing (i.e. it is the *only* available means).

On this analysis, is taking ever-such-extreme care not to injure P indeed the only available means that D has of not injuring P? One may doubt whether it meets condition (c).[3] But more

[3] That D will not encounter P at all as he weaves his drunken way home in the car is, in general, a perfectly realistic prospect. So why isn't this realistic prospect billed, in the proposed reconstruction of Brudner's argument, as an alternative available means of D's not injuring P, with the supposed result that D's taking ever-such-great care not to injure P is not left as the only available means? The answer is that it is taken for granted that the set of available means of ϕing equals the set of available means of *intentionally* ϕing. By knocking out

importantly it fails to meet condition (a). Taking ever-such-extreme care not to injure P cannot possibly be the only available means that D has not to injure P because, on closer inspection, it is not a means of D not injuring P at all.

When we read condition (a) we have to begin by grasping the relevant idea of a 'contribution' to D's ϕing. The basic idea we are looking for here is that of *making it more likely* that D ϕs. After all, for Brudner's purpose – the purpose of evaluating rival legal standards or rules – it is not enough to derive, as an object of evaluation, a one-off obligation on D to take ever-such-extreme care not to injure P as a means of not injuring P. What is needed is a *standard* or *rule* according to which D is to take ever-such-extreme care not to injure P as a means of not injuring P. So our interest is not in D's possible one-off failures and successes in not injuring P by taking ever-such-extreme care not to. Our interest is in his *prospect* of failure and success, his *tendency* to succeed when he conforms to the advertised standard or rule, the *probability* of his not injuring P in the process. Thus the acid question facing Brudner's rearguard defenders is: Is it really the case that (assuming D has the capacity not to injure P) D always continues to improve the probability of his not injuring P, the more care he takes not injure P? Is D's taking ever-such-extreme care not to injure P in *that* sense making a contribution to his not injuring P?

To that question the answer is a resounding no. In principle there always comes a point at which D's taking further care not to injure P does nothing to improve the probability of his not injuring P. Indeed, in principle there always comes a (further)

some of the competition in advance – namely all available means of *un*intentionally ϕing – this tweak of condition (c) makes it systematically easier for γing (an action performed in order to ϕ, such as 'taking care to ϕ') to be declared the only available means of ϕing. The tweak therefore makes it easier than it would otherwise be to generate a derivative obligation to γ (e.g. to take care not to injure) out of an existing obligation to ϕ (e.g. simply not to injure). The conclusion of this paper will help to show why the tweak is illegitimate.

point at which D's taking further care not to injure P is positively counterproductive. The further care D takes not to injure P, beyond this point, actually makes it more, not less, likely that D will injure P. Like other moral virtues, to put the point in familiar Aristotelian terms, the virtue of caution lies in a mean between deficiency and excess. Excessive caution, or at any rate one kind of excessive caution,[4] is the caution that fails in its own terms, the caution that tends to precipitate the very eventualities that the excessively cautious person was taking such pains to avoid. In this vein, we are all familiar with the self-defeating menaces that are the overcautious driver, the overprotective parent, and the oversolicitous lover. Set alongside their all-too-careless counterparts these characters are of course genuinely admirable.[5] But our admiration for them is tinged with pity and frustration. That is because (discounting the secondary success, the bittersweet moral victory, that lies in the mere fact of their being admirable) these too-careful characters are alas no more *successful* at what they do than their all-too-careless counterparts – and that is in spite of, nay *because* of, all the care they take to do it. Their taking such care, put simply, is not a means to their actually doing what they are taking such care to do. On the contrary it is a means to their doing the opposite. So if they have a straightforward obligation not to hurt some other person (say), that obligation yields – on any credible view of the logic of practical rationality – no derivative obligation on them to take

[4] Perhaps there are other kinds. Couldn't one be overcautious either in being self-defeatingly cautious or in being, as Brudner emphasizes, self-*effacingly* cautious? I wonder whether the latter fault is strictly speaking an instance of overcaution, as opposed to an instance of cravenness or slavishness (that being the virtue of humility taken to a self-defeating point).

[5] I do not mean that they are admirable in spite of their fault. I mean that their fault is an admirable one. It does not follow that it is any the less a fault or that the conduct exhibiting it is any the less wrong or blameworthy. On the logic of admirable yet blameworthy wrongdoing, see Michael Stocker, *Plural and Conflicting Values* (Oxford: Clarendon Press1990), ch 2.

the extreme care that they irrationally imagine they should take. Why, it does not even yield a derivative *reason* to take such extreme care. What it yields, on the contrary, is a reason, and perhaps sometimes even an obligation, *not* to take such extreme care. It is a salutory reminder of the extent of reason's cunning that a straightforward obligation not to injure P, far from yielding an obligation to take ever-so-much care not to injure P, could in principle yield a derivative obligation to do exactly the opposite, i.e. to proceed without taking very much care, or in extreme cases maybe even any care at all, not to injure P.[6]

II

I mention such dramatic and hard-to-envisage cases mainly to bring home this message. The highest level of productive care (as we might call that point beyond which taking further care makes one no more likely to do what one is taking care to do) is not a constant. It varies from activity to activity. If one must walk across a bed of hot coals it is probably best, as a rule, to be entirely careless where one places one's feet, since any care one takes (e.g. to avoid sharp coals) will, as a rule, just slow one's progress and intensify one's pain. Here the highest level of productive care is on the low side. On the other hand, if one is walking a tightrope very great care in placing one's feet is as a rule exactly what is called for, so I am told, and the highest level of productive care is very high. Thus the overcautious types we mentioned might be, in that respect, good tightrope-walkers but bad hot-coals-walkers. There are in other words very dramatic variations in where the mean of caution lies, as we shift from activity to activity. In principle, however, a highest level of productive care always lurks somewhere. Even in tightrope-

[6] The most important modern study of this phenomenon of self-defeating endeavour is Derek Parfit's *Reasons and Persons* (Oxford: Clarendon Press 1984), part 1.

walking there comes a point at which one's ever more elaborate precautions against falling (lessons in concentration, extra checks on the rope, intensive practice sessions all afternoon, etc.) tend to stop helping one walk the rope, and soon enough (thanks to the extra stress and tiredness they bring on, etc.) they become the very things that precipitate one's fall. So a straightforward obligation on D not to fall from his tightrope could not – *pace* Brudner's would-be charitable interpreters – yield a derivative obligation on D to take ever-such-extreme care not to fall from his tightrope. What it might yield is a derivative obligation on A to take all *reasonable* (i.e. productive) care not to fall from the tightrope, where this is contrasted with the *unreasonable* (i.e. excessive) care not to fall from the tightrope that would actually hasten A's falling.

When I say 'reasonable care' here, do I mean reasonable care as this is understood by lawyers – namely the care that all and only the negligent fail to take? Something close. What I have been calling the highest level of productive care is indeed one of the central determinants of what counts, in law, as the negligent pursuit of this or that activity. To be sure, it is not the only determinant. Even an extra measure of productive care is not regarded as an extra measure of reasonable care, and so is not required to avoid negligence-based liability, if it is too difficult or costly for D to lay on. But Brudner's argument plays up the difficulties and costs to D of laying on additional measures of care, without noticing that there is always, at the highest level of productive care, an independent cap on the rationality (and hence the derivative obligatoriness) of D's incurring those difficulties and costs. Because of this cap, there is no possibility of a straightforward standardized legal obligation on D not to injure P yielding a derivative standardized obligation on D to take ever-such-extreme measures of care not to injure P. If the obligation not to injure P yields any derivative obligation to take care at all, it is at most an obligation to take all *productive* measures of care – all care that contributes by making it more likely that one will

avoid what one cares to avoid – and that often means, *pace* Brudner, only very modest measures of care. In principle, as I said, it might sometimes mean no measures of care at all.

These remarks were supposed to help bring out the rational relationship between negligence-based (or fault) liability on the one hand and strict liability on the other. Essentially, the obligation that one is held to have under a regime of fault liability is an obligation *to take productive care to do* that which, under a regime of strict liability, one would merely have a straightforward unadorned obligation *to do*. As well as bringing out the connection between the two modes of liability, however, this last formulation also helps to bring out the simple ineliminable qualitative contrast between them. The contrast is but one instance of a larger contrast that turns out to be at stake in many of the deepest puzzles of moral and political philosophy. Elsewhere,[7] I have coined some terminology to mark the larger contrast. I have called it the contrast between obligations to *try* and obligations to *succeed*. An obligation to succeed is a straightforward obligation to ϕ, for the performance of which only one's actually ϕing matters, never mind what steps one takes (if any) with a view to ϕing. An obligation to try is the converse case: only the steps one takes with a view to ϕing are relevant to whether one performs the obligation, never mind whether one actually ϕs thereby. On this account, a straightforward obligation not to injure P – the obligation that D is held to have been under (and to have failed to perform) when she is held strictly liable for injuring P – is a straightforward example of an obligation to succeed. By contrast what the law calls D's 'duty of care' to P – the obligation to take care not to injure P, nonperformance of which puts D at fault and grounds D's fault liability – is a straightforward example of an obligation to try.

[7] In 'The Purity and Priority of Private Law', *U Toronto LJ* 46 (1996), 459 at 486.

Or maybe not so straightforward? Some lawyers, in my experience, are resistant to the suggestion that an obligation to take care is an obligation to try. On the one hand, it is said, no amount of trying is logically *sufficient* to perform one's obligation to take care in law, for the legal standard of negligence is an impersonal (or 'objective') standard and this means that, for at least some people, try as they might they will not perform their legal obligation. On the other hand, it is said, no amount of trying is logically *necessary* to perform one's obligation to take care in law. One takes reasonable care not to injure in the relevant sense just as long as one acts in ways that actually reduce the risk of injury one poses to an acceptable level, even if one was not at all concerned, or motivated, to reduce the risk. These objections combine to portray the legal obligation to take care as just another obligation to succeed – to be exact, an obligation to succeed in not behaving too riskily, never mind what one was trying to do – and hence they purport to eliminate the 'ineliminable qualititative contrast' that I drew between strict liability and fault liability. Both of the objections are confused, however, and my contrast stands uneliminated.

Regarding the first objection: If some people do not perform their obligation to take care however hard they try, this does not go to show that it is not an obligation to try. It only goes to show that it is an obligation to try harder (more assiduously) than they are capable of trying. True, an obligation to try harder than one is capable of trying is a problematic obligation. In particular, it runs up against what I set out as condition (b) above, namely the 'availability' condition for establishing an obligation. But recall that the availability condition, like the other conditions I mentioned, was only a condition for establishing an obligation by deriving it from another obligation via the 'only available means' argument. And recall that the suggestion that a legal obligation to take care might be an obligation derived from another obligation by the 'only available means' argument was floated only as part of an *ad hominem* bid to salvage Brudner's

attack on strict liability. If the legal obligation to take care is not after all an obligation derived from another obligation, or is derived from another obligation only by some argument other than the 'only available means' argument, then there is no reason that I can think of to assume that the obligation to take care exists only subject to availability condition – no reason to assume, in other words, that one has it only on condition that one has the capacity to perform it.[8] But even if there is reason to make this assumption regarding obligations more broadly, there is no reason to think that the assumption applies any the less to obligations to succeed than it does to obligations to try. The point is that the distinction between obligations that one has only on condition that one has the capacity to perform them and obligations (if there be any) that are not subject to this condition is a distinction that *cuts across* the distinction between obligations to try and obligations to succeed.[9] That one lacks the capacity to perform one's obligation does not entail that one has an obligation to succeed rather than an obligation to try. Rather, it remains an obligation to try if the explanation of why one lacks the capacity to perform it is that one lacks assiduousness, i.e. the capacity to try hard enough.

As for the second objection: According to this objection, negligence is a matter of the injurious tendency of one's action, and in assessing that injurious tendency intended effects (those one was trying to achieve) are regarded, through the lens of the negligence test, as being completely on a par with side-effects.

[8] That 'ought' implies 'can' is commonly given as the reason. But that is not a reason – that is a restatement of the same misguided assumption. The assumption is often ascribed to Kant but much of Kant's work was devoted to refuting it in the sense in which it is usually ascribed to him and widely but mistakenly assumed to be true. Cf note 42 below.

[9] Those who think that 'obligations to succeed that I have the capacity to perform' is inevitably an empty set are confusing the capacity to perform with the capacity to guarantee performance. Cf. Michael Moore, 'Authority, Law, and Razian Reasons', *Southern California LR* 62 (1989), 827 at 875.

All are relevant according only to their foreseeability, not according to whether or not the agent literally 'cared' (here meaning 'tried') to bring them about or avoid them. Trying assiduously enough to avoid injuring is admittedly *one* way to avoid being negligent; but one can also avoid being negligent by (fortuitously) posing only acceptable risks of injury by one's conduct. So we are often told by, for example, those who have an economistic interpretation of the law. But is this the real legal position? I think not. Let's concede – although the matter is more complicated than it looks – that the law doesn't locate D's negligence in the fact that he tried to do what he tried to do.[10] But it nevertheless clearly does locate D's negligence in what he meanwhile *didn't* try to do. The whole point is that D didn't *take sufficient care to avert* (limit, reduce, control) the injurious side-effects of his endeavours. The words 'to avert' here mean 'in order to avert', 'with a view to averting', or (in other words) 'intending to avert'. What else could be meant by them? Taking care is an essentially intentional action. One cannot take care not to ϕ without trying not to ϕ. That is also the law's position. Negligence in law is a failure to try assiduously enough to avert (limit, reduce, control) the unwelcome side-effects of one's (otherwise valuable) endeavours. It follows that the obligation that one fails to perform when one acts negligently is indeed an obligation to try. The nonperformance of an obligation to try is what gives rise to fault liability in law, just as the nonperformance of an obligation to succeed is what gives rise to strict liability. The contrast between the two is as basic as they come.

[10] The complications are brilliantly exposed in John Finnis's 'Allocating Risks and Suffering: Some Hidden Traps', *Cleveland State LR* 38 (1990), 193.

III

Brudner's failed objection to the legal imposition of strict liability was notable for having virtually turned on its head the most common objection to strict liability in the law. Brudner's objection proceeded from the thought that to avoid incurring strict liability, one needs to exercise the most extreme conceivable care. By contrast, the common objection focuses on the fact that, when strict liability is at issue, the law does not give two hoots how much care one takes. This common objection in turn comes in two variants. I will call them the *moral intelligibility* variant and the *institutional fairness* variant. According to the institutional fairness variant, it is unfair (or contrary to the rule of law, or something along those lines) to ground legal liability in nonperformance of an obligation to succeed, since this makes it impossible even in principle for people to plan their lives according to whether their actions will or will not attract legal liability. For they cannot possibly know this until they know how their actions turn out – as successes or failures – by which time it is too late to reconsider whether to perform them. Unfair? You may think so. But if you sign up to the other variant of the objection – the moral intelligibility variant – you do not even get as far as asking whether the legal enforcement of obligations to succeed is unfair. For what you deny, in a sense, is the very possibility of there being obligations to succeed, be they enforced or not. Of course, this is not to deny that the law *asserts* the existence of unperformed obligations to succeed whenever it imposes strict liability. That cannot be denied. Nor is it denied that the law may have its reasons for imposing the liability and hence for asserting the existence of the underlying obligations. Naturally it may. What is denied is that the law's assertion of the existence of the obligations can be made morally intelligible. For morally speaking – i.e. apart from the law – there is and can be no such thing as an obligation to succeed. Or so the objection

goes. Strict liability, as Thomas Nagel famously put it, 'may have its legal uses but seems irrational as a moral position.'[11]

At first sight this second variant of the common objection seems to miss the target completely. If the assertion of an obligation to succeed 'has its legal uses', why is it an objection to those legal uses that obligations to succeed could not possibly exist *apart* from the law? Why should the law care whether its assertions are morally intelligible so long as they are legally useful? The answer, of course, is that the law by its very nature claims to bind people morally; it purports to tell them what they *really* ought to do, not merely what they ought to do according to law.[12] Legal obligations, in short, are would-be moral obligations. What this means is that, if they are to be defended as legal obligations, they have to be defended as would-be moral obligations, i.e. inspected in a moral light. It does not follow, of course, that defensible legal obligations all need to be institutionalizations of moral obligations that already exist apart from the law (the nonperformance of which would be *malum in se*). Defensible legal obligations are often new moral obligations created by law (the nonperformance of which is *malum prohibitum*), and in principle the creation of these new moral obligations can be defended perfectly adequately by relying on the usefulness (the consequential advantages) of the liability regime thereby brought into being.[13] But the creation of such

[11] Thomas Nagel, 'Moral Luck', *Proc. Arist. Soc. Supp. Vol.* 50 (1976), 00; reprinted in his collection *Mortal Questions* (Cambridge: Cambridge University Press 1979), 24 at 31 (to which page references below refer).

[12] For two largely complementary defences of the thesis that such claims form part of the very nature of law, see Raz, *The Authority of Law* (Oxford: Clarendon Press 1979), chs 1 and 2, and Finnis, *Natural Law and Natural Rights* (Oxford: Clarendon Press 1980), chs 1 and 11.

[13] Thus when I say 'moral' here I do not mean 'moral' in the narrow technical sense often favoured in recent tort theory, which is by way of contrast with 'economic' or 'policy-oriented'. Contrast, for example, Ernest Weinrib's anti-*mala-prohibita* invocation of 'moral' in 'Towards a Moral Theory of Negligence Law', *Law and Philosophy* 2 (1983), 37.

legal obligations always remains subject to the following constraint: On pain of indefensibility as legal obligations, they must be obligations that could conceivably have existed apart from the law. On pain of indefensibility, it must always *make sense* to say of any legal obligation 'this is your moral obligation'. Otherwise it is impossible to judge legal obligations as would-be moral obligations. So if, as many are wont to argue, it never makes sense to say of an obligation to succeed 'this is your moral obligation,' then the assertion of obligations to succeed in law is always indefensible, whatever their 'legal uses'. Although rarely spelt out, this is the moral intelligibility variant of the common objection to the legal imposition of strict liability.

These lines of argument are invoked to discredit strict liability. But as I expressed them it seems that they also bite, in the process, against some other kinds of legal liability. Take, for example, liability for the tort of negligence at common law. In spite of its name, the tort of negligence is not the simple tort of acting negligently towards P. It is the more complex tort of injuring P *by* acting negligently towards P. To put it another way, although one has a legal obligation to take reasonable care not to injure P (the 'duty of care'), breach of this straightforward obligation to try is not a tort. It is not a legally actionable breach of obligation. What is legally actionable, and is a tort, is the breach of a more complex obligation, namely the obligation *not to injure P by* not taking reasonable care not to injure P. This is not, as you can see by the italicized words, a straightforward obligation to try. Rather it is a kind of hybrid trying-succeeding obligation. More than one kind of hybrid trying-succeeding obligation is, of course, logically possible. One kind of hybrid would be an obligation to *succeed-by-trying* – an obligation, for instance, to take reasonable care not to injure P and (thereby) to avoid injuring P. This is obviously a tall order. One would fail to perform this hybrid obligation if *either* one did not take the correct measure of care not to injure P *or* one did in fact injure P. But the hybrid obligation, nonperformance of which constitutes

the tort of negligence, is not a hybrid of this 'tall-order' kind. It is the converse case. One commits the tort of negligence only if one *both* failed to take the correct measure of care not to injure P *and* one did in fact injure P. The obligation at stake is not an obligation to succeed by trying but an obligation *not to fail for want of trying*. We might call this a 'short-order' as opposed to a 'tall-order' hybrid. It is more readily performed than either a straightforward obligation to try or a straightforward obligation to succeed, all else being equal, since there are always two alternative ways of performing it: either by trying to the requisite extent (whether successfully or not) or else by succeeding (whether or not by trying).

The fact that it is more readily performed than either a straightforward obligation to try or a straightforward obligation to succeed does not entail, however, that the presence of such a hybrid obligation in the law is also easier to defend. True, one can quickly nuance the institutional fairness argument so that it cuts against the legal enforcement of straightforward obligations to succeed (i.e. strict liability) but not against the legal enforcement of short-order hybrids (as in the action for negligence at common law). The institutional fairness argument is indeed most often ventured asymmetrically, with the focus on helping those who want to be sure in advance that they will *not* incur any liability, rather than on helping the perhaps more eccentric types who want to be sure in advance that they will incur it. Since taking the legally specified measure of care is logically *sufficient* to extinguish liability for the tort of negligence, one can plan to steer clear of that liability just as readily as one could plan to steer clear of liability for non-performance of a simple obligation to take the specified measure of care (i.e. a straightforward obligation to try). True, taking a legally specified measure of care is not logically *necessary* to extinguish liability for the tort of negligence, since that liability could also be extinguished by the happy fortuity that one injured nobody by one's carelessness. But naturally this possibility is no skin off the

nose of people who merely want to plan safely to avoid the liability, for they only need to know what investment of care would be sufficient to do so. If these people are the only ones that the institutional fairness argument seeks to protect, then that argument cuts against strict liability (the legal enforcement of straightforward obligations to succeed) but not against liability for the tort of negligence (the legal enforcement of short-order hybrid obligations not to fail for want of trying).

Matters are not so clear-cut, however, when we come to the moral intelligibility objection. If straightforward obligations to succeed (and hence strict liability torts) turn out to be morally unintelligible, could obligations not to fail for want of trying (and hence the tort of negligence at common law) nevertheless turn out to be morally intelligible?[14] The final answer to this question depends, of course, on what exactly it is that is said to *make* obligations to succeed morally unintelligible. For the time being, however, our interest is not in the answer. Our interest is in the question itself. As followers of modern policy debates about the law of torts, especially those that rage in the North American law schools, we are all accustomed to think of the history of modern tort policy as a struggle for supremacy between the reigning tort of negligence on the one hand, and, on the other hand, the upstart strict liability torts (on the *Rylands* v *Fletcher*[15] model) that occasionally surface, or are promoted, as pretenders to its throne (nowadays, for example, in the field of products liability). In the mentality of modern tort law and tort theory, to defend the one is implicitly to attack the other. But our question puts a completely different complexion on the relationship between them. It shows that in another dimension the tort of negligence may on the contrary have its moral fate *bound up* with that of its strict liability rivals. For the tort of negligence and its strict liability rivals on the *Rylands* v *Fletcher* model stand together in

[14] See Nagel, 'Moral Luck', above note 19, 28-29.
[15] (1866) LR 1 Ex 265; [1868] LR 3 HL 330.

opposition to another little-noticed but increasingly powerful pretender, a model of tort liability in the incurring of which it matters not one jot whether D actually injured P, but *only* how much care D took that P not be injured. The really stark contrast is, of course, the contrast between this kind of 'bare negligence' tort (which consists in nonperformance of a straightforward obligation to try, never mind one's success) and a strict liability tort (which consists in nonperformance of a straightforward obligation to succeed, never mind whether or to what extent one tried). In this contrast, the tort of negligence at common law is logically non-aligned. It consists in the nonperformance of a hybrid trying-succeeding obligation (to be exact, an obligation not to fail for want of trying), and hence shares as much and as little of its logical structure with strict liability torts as it does with bare negligence torts.

But is it also *morally* non-aligned? There are two dramatically contrasting ways of thinking about the moral intelligibility, and hence the defensibility, of the tort of negligence at common law. It is possible to take compromise positions intermediate between them but they represent the two magnetic poles of the debate, the first of which has lately proved more magnetic than the second. The first interpretation – the one which has come into the ascendant in the late-twentieth-century literature – would have it that the tort of negligence is essentially a variation on the theme of a bare negligence tort. The essence of D's wrong, morally speaking, lies in his failure to take adequate care not to injure P. The extra condition that P actually have been injured by D's failure to take care is regarded, correspondingly, as morally secondary. Its defence is accordingly apt to be a parasitic defence, i.e. one that presupposes that the real moral wrong has already been done to P. Thus the extra condition is collaterally defended as, for example, a *locus standi* condition calculated to

optimise the incidence of litigation in respect of the wrong,[16] or as a way of ensuring that compensation for the wrong will in principle be quantifiable without resort to excessive speculation about probabilities. But by the same token the door is left open to allowing the extra condition that P be injured by D to be eliminated or watered down while leaving the central moral theme of the tort in place. On this interpretation, one can imagine modifying the tort such that in some cases (e.g. mass injuries arising from multi-manufacturer pharmaceuticals or industry-wide environmental hazards) it is no longer a condition of D's liability for the tort of negligence that D actually injured P, but only (say) that P was injured and D was one of a finite group of careless operators who may well have injured her, or even (say) that P was not injured at all but only left insufficiently protected against injury by D and his fellow operators. These are regarded, on the first of the two contrasting interpretations that I have in mind, as relatively marginal adjustments to a tort, the moral essence of which lies in nonperformance of a straightforward obligation to try, viz. the common law's famous duty of care.

But on the other interpretation that I have in mind these adjustments are far from marginal. They represent the most dramatic possible abrogation of the moral foundation of the tort of negligence. For in spite of the tort's misleading name, the moral essence of D's tort in a negligence case is really just that he injured P. In other words, the tort of negligence at common law is morally speaking a variation on the strict liability model of a tort, in which what is of the essence is what one actually does (injures P), never mind what one merely tries to do (one's not taking adequate care not to injure P). The extra condition that D not have taken adequate care not to injure P is regarded, correspondingly, as morally secondary. Its defence is accordingly

[16] See, most obviously, Richard Posner's influential 'A Theory of Negligence' *J Leg Stud* 1 (1972) 29 at 46-52.

apt to be a parasitic defence, i.e. one that presupposes that the real moral wrong has already been done to P. The negligence condition is defended, for example, as one that is needed merely to meet the institutional fairness objection. The moral wrong one does is essentially just to injure P, but it is unfair to hold D liable for that wrong without providing a way in which D could have taken steps to guard himself against incurring the liability. Thus the extra condition of fault is grafted on to the tort for institutional reasons. Or some such parasitic, morally secondary reasons. Correspondingly the negligence condition could in principle be removed in appropriate cases without eating away at the moral foundations of the tort. Where someone embarks on ultra-hazardous activities, for example, they may be regarded as having been put on fair warning of their liability for the injuries they do merely by the fact that their activities were ultra-hazardous. They could plan not to incur the liability simply enough, by giving up the blasting or the chemical processing that is creating the hazard. So under some conditions – e.g. those prevailing in *Rylands* v *Fletcher* – the extra negligence condition may defensibly be dropped, and then we are left with the moral essence of the tort, which emerges as the simple unadorned wrong of *D injuring P*.[17] Or so goes the rival – and nowadays markedly contrarian – interpretation of the tort of negligence.

[17] We can see here one reason why Brudner may have thought that strict tort liability demands of D that he be entirely self-effacing. It is not because tort liability is strict but because in modern legal systems it is typically strict *and conditional*, i.e. it is a strict liability that arises only when one is engaged in certain pursuits, such as blasting and manufacturing consumer products. These extra conditions are needed to meet the problem of institutional fairness. To guarantee avoiding the strict liability, it is logically sufficient (but notice: not necessary) that one give up the pursuit in question. Now surely *that* would be a huge (self-effacing) burden for one to bear? True, but one cannot plead this as an objection to the liability's being strict. Suppose, in response, the law called one's bluff. Suppose it removed the condition and left one with an unconditional strict liability instead, applicable as much to car-driving and hairdressing as to blasting and consumer manufacturing. Then the law would

IV

We owe to Tony Honoré the most important modern attempt to rehabilitate this contrarian interpretation of the tort of negligence at common law.[18] When we put the tort of negligence side by side with modern strict liability torts, argues Honoré, the first feature that calls for our defensive or critical attention is not the much-discussed feature that divides the two ('fault' v 'no-fault'), but the less-widely-remarked-upon feature that they both have in common, namely what Honoré calls their shared element of *outcome responsibility*. Depending as it does on the moral defensibility of outcome responsibility, the tort of negligence is in the same boat, morally speaking, as a strict liability tort. The tort of negligence is essentially a variation on the theme of a strict liability tort, in which the basic wrong lies in D's actually injuring P. Accordingly, those who want to defend the tort of negligence at common law had better begin by defending, not attacking, the morality of strict liability. There is plenty of time to part company with enthusiasts for strict liability later, when one comes to the question of what *further* conditions must be met before legal liability is justified. In answer to that question, enthusiasts for strict liability may say 'anything but a fault condition' while enthusiasts for the tort of negligence may say 'a fault condition'. That polarization is for later. Before that, according to Honoré, comes the morally more basic problem of

have met one's objection head on by removing one's reason to give up the blasting in favour of, say, hairdressing. Yet rather than eliminating strict liability the law's reaction would have expanded it to take in hairdressing as well. The point I am making is that if strict liability in modern tort law seems to require that one efface oneself by taking extreme measures to avoid liability, that is usually because there is not enough of it around to make the self-effacement utterly fruitless.

[18] Largely in the papers collected in his *Responsibility and Fault* (1999), although anticipated in the preface to the second edition of Hart and Honoré, *Causation in the Law* (1984). In the following footnotes, all references to papers reprinted in *Responsibility and Fault* use the pagination of the reprint.

understanding the common ground that unites the two competing sets of enthusiasts against everyone else, namely the moral significance that they agree in attaching to *D's actually injuring P* – the element of outcome responsibility.

It should be said that Honoré's first articulation of this contrarian line of thought suffered from an equivocation. At first he made the wrong connection between strict liability and liability for the tort of negligence. He said that it was the *objective standard of care* in the tort of negligence which gave that tort an affinity with strict liability torts. 'For the objective standard of competence,' he wrote in his first major venture into the subject, 'imposes a form of strict liability on that minority of shortcomers who cannot achieve it.'[19] We have already seen that this is a mistake. To defend strict liability is to defend, minimally, the moral intelligibility of straightforward obligations to succeed. An obligation to try harder than one is capable of trying, of the kind that is created when 'shortcomers' encounter the objective standard of care, is not an obligation to succeed. It remains an obligation to try. As we saw already, the distinction between obligations to try and obligations to succeed *cuts across* the distinction between obligations that one has only to the extent that one has the capacity to perform them and obligations (if there be any) that are not subject to this condition. The problem of strict liability correspondingly cuts across, rather than tracking, the problem of the objective standard of care in the tort of negligence. By lumping the two together Honoré equivocated about which of the two problems he was really trying to tackle.

Honoré was led astray here by his instinct to interpret the common objection to strict liability as an objection to the intelligibility (or, on its other variant, the fairness) of the law's exposing people to 'moral luck', i.e. luck in whether they come up to scratch in what they do and hence in what judgments they

[19] 'Responsibility and Luck', *Law Quarterly Review* 104 (1988), 530, reprinted in *Responsibility and Fault*, 14 at 22.

are properly open to. Surely, he reasoned, not only strict liability but also the objective standard of negligence can be objected to on this ground? True enough – if the objection makes sense. But, as Nagel established, it does not. There can be no such thing as a coherent general objection to our being exposed to moral luck. Attempts to explicate such an objection are an object lesson in the hazards of argumentative overkill. For what counts as luck is always, Nagel shows, luck *only relative to some baseline or other.* Whenever something is held to be luck, there is necessarily something else that is held *not* to be luck, and it is only relative to this second thing that the first counts as lucky or unlucky. The problem with a general objection to our exposure to moral luck is that *everything* we do is entirely a matter of luck relative to *some* baseline or other. Even when I maliciously attempt to do away with my professional rival – on any credible view an action that is open to some moral judgment – my doing so is luck relative to some baselines (e.g. the baseline of my genetic make-up, which pre-programmed no such behaviour and was compatible with my leading a life without any such attempt). It follows that to object to moral luck *tout court* is to object to morality *tout court.* Indeed it is to object to judging people's actions by any standards at all. As Nagel himself puts it: 'The area of genuine agency, and therefore of legitimate moral judgment, seems to shrink under this scrutiny to an extensionless point.'[20]

Misreading Nagel's tone here, some took it that they were supposed to endorse the conclusion of this sentence, and to regard morality itself as having been discredited by Nagel's demonstration that there is no place for morality to hide from luck. But in fact Nagel's argument is a classic *reductio.* The conclusion is absurd – agency *does* have some reach and moral judgment *does* have some area of application – so something must

[20] 'Moral Luck', above note 19, 35.

have gone wrong in the premisses.[21] What went wrong, Nagel showed, was the instinct to formulate as a general objection to the intelligibility (or fairness) of exposing people to 'moral luck' what were really different people's quite different and unrelated objections to the use of different baselines as the baselines for our moral judgments. Some people think that the question of whether we come up to scratch in what we do should be relativised (baselined) to our own personal capacities, and in the light of that view they rebel at objective standards of care as used in the law. Some think – an overlapping constituency, I suspect – that the question of whether we come up to scratch in what we do should be relativised (baselined) to the possible limits of trying, and in the light of that view they rebel at strict liability. Each of these suggested baselines needs to be defended, and for that matter objected to, on its own merits. Maybe refusing to acknowledge one or other or both of these baselines as fixing the proper boundaries of moral judgment is morally unintelligible or institutionally unfair. But the observation that in refusing to acknowledge them we expose people to 'moral luck' does absolutely nothing to tell us *what it is* that would make such reliance either unintelligible or unfair. For (as Nagel shows) we are automatically exposed to moral luck, against some baseline, just in virtue of the fact that our actions are held up to judgment at all. The so-called 'problem of moral luck' (like the so-called 'problem of free will' with which it is supposedly connected) is therefore a pseudo-problem.[22] It represents the bundling

[21] Even if you don't accept the assertions, Nagel and Honoré do. So maybe I should say that the argument is a classic *reductio* for their purposes.

[22] An irony: Bernard Williams and Tom Nagel conjured up 'the problem of moral luck' as a topic for their joint Aristotelian Society seminar in 1976. Both of them showed, in their different ways, that there is no such problem. Yet strangely the problem took off while its agreed dissolution was forgotten. Williams' eponymous paper appears before Nagel's in the original periodical printing (above note 19), and is reprinted in Williams, *Moral Luck* (Cambridge: Cambridge University Press 1981).

together of various different problems about the grounds of moral judgment that have nothing at all in common save that they are all problems about the grounds of moral judgment. Two of these problems – and they are as distinct from each other as it is possible for two problems to be – are the problem of the objective standard of care (can there be moral obligations to do more than one is capable of doing – whether in the way of succeeding or trying – and if so should they be legally enforced?) and the problem of strict liability (can there be straightforward moral obligations to succeed – whether within or beyond the limits of one's capacities – and if so should they be legally enforced?)

Fortunately it did not take Honoré long to recover from the effects of this distracting conflation. A few pages later he puts it behind him and trenchantly captures the real respect in which the tort of negligence is morally aligned with strict liability torts of the *Rylands* v *Fletcher* variety, and to that extent stands or falls with them. It turns out to have nothing to do with the objective standard of care. Rather, explains Honoré,

> [s]trict liability is one species of enhanced responsibility for outcomes. This does not entail that whenever a harmful outcome is properly allocated to someone, this justifies imposing on him a strict liability to compensate for that outcome. ... [R]esponsibility for a harmful outcome should not automatically involve a legal duty to compensate. An extra element is needed to ground the legal sanction. Sometimes [as in the tort of negligence at common law] the extra element is fault. ... For strict liability [as under *Rylands* v *Fletcher*] the extra element is usually that the conduct of the harm doer carries a special risk of harm of the sort that has in fact come about.[23]

There admittedly comes a point at which one needs to decide what further conditions, if any, one will insist upon for legal liability: maybe (objective) fault, as in the tort of negligence, maybe not, as under *Rylands* v *Fletcher*. But that little internecine

[23] 'Responsibility and Luck', above note 19, 27.

squabble is for later. First, in the order of moral argument, comes the question of what kind of standard these further conditions of liability are supposed to be grafted onto, and why. In the tort of negligence and *Rylands* v *Fletcher* torts alike, argues Honoré, the underlying standard is a standard of outcome responsibility, a standard which attaches a person (D) to the way his actions actually turn out (say, with P's being injured).

I think Honoré chooses the label 'outcome responsibility' to designate the issue that he is interested in here because the label's intriguing ambiguities correspond to some intriguing ambiguities in the issue itself, and these Honoré understandably wants to keep alive for investigation. For a start, talk of an action's 'outcome' equivocates nicely between a reference to an action's *consequences* and a reference to (what some philosophers call) its *results*.[24] The consequences of an action are what they sound like: they are eventualities that follow an action and that are also causally connected to it. Results are different. They are the causal *constituents* of actions, i.e. they do not follow the action but form part of it.[25] P's death, for example, is a consequence of D's trying to kill P. On the other hand P's death is a result of D's actually killing P, because the action of killing P (unlike the action of trying to kill P) is partly constituted by P's dying. Until P is dead, D hasn't killed him but has only tried to. Death is a consequence of the one action and a result of the other, but it is perfectly natural to say, in the intentionally ambiguous terminology chosen by Honoré, that it is the 'outcome' of them both.

[24] This useful terminology is G.H. von Wright's, from *Norm and Action* (London: Routledge and Kegan Paul 1963), 39ff. The distinction is interestingly nuanced in Anthony Kenny, *Will, Freedom and Power* (Oxford: Basil Blackwell 1975), 54ff.

[25] I call a result a 'causal' constituent because the action of which it forms part is an action of causing that result (or occasioning it, or provoking it, or inducing it, or standing in some other causal relationship to it).

As for 'responsibility', this is a notoriously multivocal term.[26] For Honoré's purpose, the interesting ambiguity that it harbours is this one. In some contexts 'my responsibility' means something very close to 'my liability'. Being responsible in this sense means facing the adverse normative consequences of having failed to perform one's obligations (such as being subject to a power of punishment or being subject to a new obligation to pay compensation or apologise). But on other occasions 'my responsibility' refers to something else: it refers to the obligations themselves rather than to the normative consequences of their nonperformance. Thus failing in one's legal responsibility (=obligation) may have, as one of its legal consequences, legal responsibility (=liability). Actually, to be strict, a responsibility in the 'obligation' sense is not exactly the same as an obligation, so the '=' in the first set of parentheses here is slightly misleading. Rather a responsibility in this sense, like a right, is the *ground* of (one or more) obligations.[27] Often, indeed, D's obligations to P are grounded in the combination of P's rights and D's responsibilities. But just as we did with rights, we may put this subtlety on one side for present purposes. That is because just as there is nothing that counts as D violating P's rights other than D failing to perform her obligations to P, so there is nothing that counts as D failing to fulfil her responsibilities to P other than D failing to perform the associated obligations. Thus if we are interested in the ingredients of D's *wrongs* – for instance, her torts or breaches of contract – the difference between an obligation and the responsibility that grounds it need not concern us. Nor, for that matter, need we concern ourselves with the similarly tangential conceptual wedge that can be driven between liability

[26] The best study of the various concepts of responsibility and the relationships among them remains H.L.A. Hart's 'Varieties of Responsibility', *Law Quarterly Review* 83 (1967), 346.

[27] This explains why, as Hart notices but does not satisfactorily explain (ibid. at 347), some obligations are not associated with any responsibilities, while others are associated with more than one. The same goes for rights.

and responsibility-in-the-liability-sense. We have grasped the pertinent aspects of the two concepts of responsibility that bear on Honoré's discussion if we think of them as responsibility (≈liability) and responsibility (≈obligation).

These twin ambiguities – outcome (=consequence? =result?) and responsibility (≈liability? ≈obligation?) – combine nicely to provide Honoré with the ambiguity in the expression 'outcome responsibility' that he wants to preserve for the purposes of his discussion. One of the issues that he is interested in under this heading is the issue of whether it is *fair* to hold people responsible (≈ liable), whether in law or in other social settings (e.g. games, professional codes of practice, traditional codes of honour, etc.), for the actual (as opposed to intended or projected) outcomes (=consequences *or* results) of their actions. The other issue he is interested in under the 'outcome responsibility' heading is the issue of whether we can really *have* responsibilities (≈obligations) to perform actions defined in terms of their outcomes (=results). It does not take long to see that these two issues correspond to the two variants I mentioned of the common objection to strict liability. The first corresponds to the institutional fairness variant, which was an objection to (legal or other) liability based on non-performance of obligations to succeed. Honoré notices that if the institutional fairness variant of the objection is successful, it bites not only against forms of liability that depend on what one did-including-results (i.e. one's success) but likewise against forms of liability that depend on what happened when one did it (the actual consequences of one's actions). Thus he rightly reads the institutional fairness objection as an objection to outcome (=results *or* consequences) responsibility (≈liability). The other issue he is interested in is, however, the deeper one of the two. It is the issue of whether it can ever be one's responsibility (≈obligation) to perform actions that have a certain outcome (= result). Only if it can be do we ever get to the further question of the fairness of making us legally (or otherwise) liable when we fail to perform such obligations. This deeper issue is the very

same one that is raised by those who deny the moral intelligibility of obligations to succeed. To deny the moral intelligibility of obligations to succeed is to deny the very possibility of outcome responsibility in Honoré's second sense, never mind its fairness. It is to deny that our responsibilities (≈obligations) really do extend beyond our merely trying to do things so as to include our actually succeeding in doing them (i.e. acting with specified results). Hence it is to deny that we can make morally intelligible a law that asserts that our responsibilities do so extend.

Corresponding to these two different explanations of what outcome responsibility is,[28] Honoré offers two quite different defences of it. One, designed to meet the institutional fairness objection, is his well-known and carefully crafted 'betting' argument. Action is by its nature a gamble, he says: some you win, some you lose. Facing liability when you lose is only half of the story of outcome responsibility; the other half includes all the *positive* normative consequences (admittedly mostly laid on extra-legally) that flow from doing what you do successfully. So long as the mechanism really does cut both ways in social life taken as a whole, one has no complaint of unfairness merely because the law concentrates on the down side.[29] Although I agree that this

[28] Elsewhere in this volume, Stephen Perry calls them respectively the 'social' and the 'personhood' senses of outcome responsibility.

[29] 'Responsibility and Luck', above note 19, 24-29. Let me mention just one point about the 'betting' argument in passing, because it supplements some remarks I have already made about the structure of Honoré's enterprise. In pursuing the 'betting' argument, it seems to me, Honoré gradually slips back into his initial mistake of confusing objections to strict liability with objections to the objective standard of care. As the 'betting' argument proceeds he starts to engage with an imaginary objector who says: some lose more than they win and some win more than they lose. How is that fair? Honoré responds with an (implausibly) optimistic view of ordinary people's failure rates that is supposed to marginalize the problem of net losers (at 28). He should have responded much more robustly by saying that the new imaginary objector was changing the subject and alleging a totally different unfairness. The new imaginary

argument is sufficient to its task, I also tend to think that it is unnecessary. In my view, one can overcome the institutional fairness objection to (at any rate some pockets of) strict liability in the law with far less elaborate manœuvres that leave fewer hostages to fortune.[30] But I am not going to pursue this matter in what remains of this paper. Instead I am going to focus my attention on Honoré's second argument.

This second argument is offered not as a rebuttal of the institutional fairness objection but rather as a response to the moral intelligibility objection. Can one even make sense of the idea that people have obligations to succeed? Can one make sense of the idea, in other words, that they are outcome-responsible in the second of Honoré's two senses? Not only can one make sense of this idea, answers Honoré. More to the point, *one cannot manage without it*. It is not the presence of outcome responsibility that makes no sense, but rather its absence:

> [O]utcome allocation can be defended on grounds deeper than the overall balance of benefit over burden; and so, in its wake, can strict

objector was objecting to the unfairness of the law's failure to relativize legal liability to people's varying capacities. Honoré's 'betting' argument was not, however, a defence of the law's failure to relativize to capacities. It was a defence of the law's refusal to ignore the importance of success as well as the importance of trying. All that one can ask of the betting argument is that it eliminate the unfairness that it was devised to eliminate, viz. the unfairness of outcome responsibility. That it leaves another supposed unfairness untouched is not, so to speak, its problem. I say 'supposed' unfairness because I believe that, barring special cases, there is nothing unfair about people being held up to standards that personally they are unable to meet. The robust (and correct) answer to this complaint of unfairness is that, barring special cases (e.g. young infants and the seriously mentally ill), people *should* be able to meet the standards in question and have no complaint if they are judged by them when they cannot.

[30] In my view one overcomes the objections simply by putting potential Ds on fair warning that they are embarking on an activity (e.g. blasting) in respect of which strict liability applies to them. Cf. note 17 above on the 'ultra-hazardous activities' condition as a fair warning condition.

liability. For outcome allocation is crucial to our identity as persons; and unless we were persons who possessed an identity, the question of whether it was fair to subject us to responsibility could not arise. If actions and outcomes were not ascribed to us on the basis of our bodily movements and their mental accompaniments, we could have no continuing history or character. There would indeed be bodies, and associated with them minds. Each would possess a certain continuity. They cold be labelled A, B, C. But having decided nothing and done nothing these entities would hardly be people.[31]

This approach to the problem is exciting and unusual. The approach is to fight fire with fire. If the absence of outcome responsibility really makes no sense then its presence can't *but* make sense, and those who think it unintelligible must therefore be thinking fallaciously. Honoré's remarks therefore hold out the tantalizing promise of a philosophical role reversal, with 'moral intelligibility' objectors to the legal enforcement of obligations to succeed finding themselves, for once, in the defensive position, with their own challenge thrown back at their feet. Nor is the challenger content with rescuing obligations to succeed from oblivion; he also aims to elevate them, in the process, to a kind of moral priority. Remember Brudner's view that a regime of strict liability takes us to the point of total self-effacement, to the point at which, as agents, we are obliterated from the world? On the contrary, according to Honoré. A regime of strict liability represents the starkest possible *reaffirmation* of our agency and its importance in the world, because the simple idea at the heart of strict liability – the idea of outcome responsibility – is the idea that we leave traces of ourselves forever imprinted on history, in the form of the countless welcome and unwelcome events that were (as Honoré puts it elsewhere) 'unequivocally our doing'.[32] In what we merely try to do this imprint is lacking, and the power of our agency is therefore but meanly represented. In that sense, trying to ϕ is secondary. Success – actually ϕing – is

[31] 'Responsibility and Luck', above note 19, 29.
[32] *Causation in the Law*, above note 18, at lxxxi.

primary. 'It is only this primary outcome responsibility,' observes Honoré, 'that can explain why we (rightly) judge murder more serverely than attempted murder, and causing death by dangerous driving more severely than dangerous driving.'[33] And it is only this primary outcome responsibility – he could have added – that puts the requisite moral distance between someone who merely fails to take an adequate measure of care not to injure P, but fortunately leaves P uninjured, and someone who fails to take that same adequate measure of care, by the same margin, *and actually injures P in the process*. For first and foremost, the deeper argument goes, we are what we do – complete with results.

V

This deeper argument of Honoré's is never developed to the same level of specificity as his 'betting' argument. In subsequent writings, indeed, Honoré has preserved the speculative and exploratory tone of the words just quoted.[34] It is the tone of a philosophical promissory note. For this reason the force of Honoré's remarks – the decisive argument that they tantalizingly promise – has never been fully brought home to those who doubt the moral possibility, let alone the moral priority, of straightforward obligations to succeed. Yet Honoré's remarks do alert us to the basic steps of just such a decisive argument. True, the argument does not quite establish the moral intelligibility of straightforward obligations to succeed, let alone their moral priority over obligations to try. What it does establish, when fully spelt out, is the moral intelligibility of *reasons* to succeed, and one important sense in which those reasons have moral priority over

[33] 'Responsibility and Luck', above note 19, 31.

[34] For instance, in 'The Morality of Tort Law: Questions and Answers' in D.G. Owen (ed), *Philosophical Foundations of Tort Law* (1995), also reprinted as ch 4 of Honoré, *Responsibility and Fault*, above note 18, at 76–77.

mere *reasons* to try. With this much established, the most important source of doubt about the possibility of straightforward obligations to succeed – namely skepticism about the independent rational salience of success – is roundly despatched. Of course, at this point the question may still in principle be raised of whether these reasons to succeed with their special moral priority can ever be *categorical mandatory* reasons to succeed – that is to say, obligations.[35] But once we see that there are indeed moral reasons to succeed and that they are in one way more basic than moral reasons to try, the available arguments for doubting that they can be categorical and/or mandatory reasons are few and unconvincing. Maybe more convincing ones could be found. But the challenge to find more convincing ones is at any rate decisively thrown at the feet of the doubters, as Honoré promised us that it would be.

To see the real strength of Honoré's argument, one needs to focus on nothing so much as its apparent weakness. 'Having decided nothing and done nothing,' concludes Honoré briskly, 'these entities would hardly be people.' Surely this conclusion goes much too far, much too fast? In the first place, for the purpose of telling the story of our lives we can surely insist on the significance of the events out in the world that were 'our doing', without allowing that this significance was necessarily a moral significance? Can't things be part of the story of what we did, in an autobiographically pertinent sense, without being part of the story of our *right*doing and our *wrong*doing? Couldn't what we do *full stop* include our successes (and failures) and yet what we do *wrong* or what we do *qua moral agents* only extend as far as our attempts (and neglects)?[36] But even if we postpone this first

[35] On obligations as categorical mandatory reasons, see John Gardner and Timothy Macklem, 'Reasons, Reasoning, Reasonableness', in Jules Coleman and Scott Shapiro (eds*) The Oxford Handbook of Jursiprudence* (Oxford: Oxford University Press 2001, forthcoming)

[36] This is the gist of the response to Honoré proposed by Arthur Ripstein in his contribution to this volume. Ripstein distinguishes the undoubted 'first-

worry for a moment – which is exactly what I propose to do – there is a second, and possibly more alarming, one. Honoré seems to be claiming that if one excises results from the story of our rightdoing and wrongdoing, it is not only the case that there was nothing we did (in the sense of nothing out in the world that was our doing) but also nothing we *decided*. Surely, on the contrary, what we *decided* is among the things that are left over when we excise from the scope of our moral agency those events in the world that were our doing in Honoré's sense? Deciding, to put it another way, belongs to the 'trying' side of the trying/succeeding divide. That reminds us that there is after all a story left over when results are excised from the story of our moral agency. It is the story of our trying, a story of our endeavour rather than our achievement. And our deciding belongs to that very story. So how can it be said, as Honoré says, that in the absence of outcome responsibility we would have *decided* nothing as well as done nothing, so that there would be *no* story left of us as human agents, as opposed to a *different* one?

Let me explain how exactly it can be said. Deciding, and for that matter trying, are actions of a logically parasitic type. One does not merely decide full stop or try full stop. Necessarily, one decides to ϕ or tries to ϕ, where ϕing is another action. So necessarily there exists, whenever one tries or decides, some further action ϕ such that one decides or tries, as the case may be, to perform it. What is more, the kind of parasitism involved here is a distinctively *rational* kind of parasitism. The 'to' in the expressions 'trying to' and 'deciding to' (like that in 'taking care to') is the familiar intention-implicating 'to' that we also find in the expressions 'with a view to', 'in order to', and 'intending to'. To be exact, trying to ϕ is acting with a view to ϕing, while deciding to ϕ is (one way of) preparing oneself, with a view to ϕing, to act with a view to ϕing (roughly, it is trying to make it

person' importance of outcomes (results, successes, achievements) from their possible 'third person' (or moral) irrelevance.

the case that one will try to ϕ). In these characterizations, acting 'with a view to ϕing' means, in turn, acting for the reason (*inter alia*) that one's action will (supposedly) contribute to one's ϕing. That one's action will contribute to one's ϕing is a reason for that action, however, only if one also has a reason to ϕ. Thus to act for the reason that one's action will (supposedly) contribute to one's ϕing is possible only if one is prepared to regard or treat oneself, at least for present purposes, as having a reason to ϕ. This is not an optional extra. Someone who really thinks that they have no possible reason not to injure P – i.e. no possible reason to *succeed* in not injuring P – cannot conceivably try not to injure P or decide not to injure P, for they cannot conceivably act or prepare to act for the reason that what they do will not injure P. Accordingly, if it is impossible to make sense of the idea of a reason to ϕ – where ϕ signifies the successful action which one is trying or deciding to perform – it is also impossible to make sense of the idea of a reason to try to ϕ or to decide to ϕ. Honoré is right, then, to think that if our success turns out to be rationally insignificant across the board, then our trying and deciding (etc.) cannot but be rationally insignificant too. Assuming, then, that his conclusion is about rational significance – about what belongs to the story of our lives *as rational agents* – he is spot on. To deny that success can have independent rational significance is to leave us without *any* story of our lives as practical reasoners. It is not merely to leave us with a story of our lives as practical reasoners that omits our successes and failures but includes our attempts, decisions, precautions, neglects, etc. For the latter story depends for its intelligibility on our granting the intelligibility of a more complete – or dare I say morally richer? – story in which our successes and failures are registered as rationally significant too.

An unargued assumption of many moral philosophers as well as almost all economists and decision theorists – and I suspect also of most lawyers – is that a reason to ϕ is the very same thing as a reason to try to ϕ. Indeed the statement 'D has a reason to ϕ' is often promptly interpreted, without explanation, as 'D has a

reason to try to ϕ.' But the above remarks already show that this is a serious mistake. A reason to try to ϕ is not a reason to ϕ but a reason to act for the reason (*inter alia*) that one's so acting will (supposedly) contribute to one's ϕing. This difference makes possible a variety of important asymmetries between one's reasons to ϕ and one's reasons to try to ϕ. I will mention three.

(i) Where acting with a view to ϕing would not contribute to one's ϕing, one has reasons to ϕ (i.e. reasons to succeed) without corresponding reasons to try. I am on a clifftop miles from anywhere looking down helplessly on a man drowning in the stormy sea below.[37] Because no amount of trying would bring me closer to success, my reasons to save the man do not yield any reasons to try to save him. Yet I still have the same reasons to save him that I would have if doing so were perfectly straightforward. That is why the situation is so horrifying. If my reasons to save him were eliminated by the impossibility of my doing so then my not saving him would be nothing to me; it would leave no trace on my life as an agent; I could walk away without compunction. As it is I am merely blocked from doing as these reasons would have me do because there is no contributory action – including trying to save – that I have any (derivative) reasons to perform. The situation is an unusually stark variant of a common one in which one has more reason to succeed than one has (derivative) reason to try, thanks to the fact that one's prospects of succeeding by trying are limited.[38]

(ii) I add 'derivative' in parentheses here because I do not mean to deny that one may have *additional* reasons to try that do not

[37] For more on this case see my 'Justifications and Reasons', in A.T.H. Smith and A.P. Simester (eds), *Harm and Culpability* (Oxford: Clarendon Press 1996).
[38] I mean one's prospects of succeeding by any amount of trying. The problem under scrutiny here is different from the problem of self-defeating endeavour encountered earlier, in which the problem was that of how *much* one tried, granting that one had every reason to try.

correspond to reasons to succeed. This is an asymmetry in the opposite direction. Just as one may have less reason to try than to succeed, so one may have more reason to try than to succeed. The extra reasons to try in such a case are non-derivative reasons to try: they are reasons to try quite apart from the contribution that trying makes to success. Suppose that it is my daughter or my brother drowning in the sea below. Then I have additional reasons to try to save her or him, since such an attempt may be an expression of love even if doomed (or perhaps: even more of an expression of love *because* it is doomed). Naturally it is possible to act for these additional reasons, and possibly, in some cases, one does not express what those reasons would have one express by one's doomed attempt unless one actually acts for them. But it is not possible to act for these additional reasons *alone*, because unless one *also* aimed at success (i.e. also acted for the derivative reason that one's trying would supposedly contribute to one's succeeding) one just wouldn't be trying – and so, obviously, one wouldn't be doing what one's non-derivative expressive reasons to try would have one do. This was why I included the parenthetical '*inter alia*' in my characterisation of reasons to try: A reason to try to ϕ is not a reason to ϕ but a reason to act for the reason (*inter alia*) that one's so acting will (supposedly) contribute to one's ϕing. There are other reasons to try apart from those based on the contribution that trying makes to succeeding, and perhaps some of these other reasons to try are obligations. One can try *for* these additional reasons. But the (supposed) contribution that one's trying makes to succeeding must be among one's reasons for trying *or else one just isn't trying*.

(iii) Does it follow that there are no possible reasons to try unless it is the case that there is at least one reason to succeed, and that trying would contribute to one's success? You may think that this follows, but it doesn't quite. What follows from the foregoing is that, to have any reason to try, the agent must *believe* or *accept* that she has a reason to succeed and must *believe* or *accept*

that trying would contribute to succeeding. That was why I included the parenthetical 'supposedly' in my characterisation of reasons to try. I could try to save the drowning man if I mistakenly thought that I could just possibly save him and that I had a reason to do so. In this case, I could have reasons to try to save him. Of course I don't have the particular reason that I take myself to have, viz. the derivative reason. About that I am mistaken. But I might still have the non-derivative reasons, e.g. that of expressing my love for him, and I can do as these reasons would have me do by acting for what I mistakenly take to be the derivative reason.

These are just a few of the many complexities that figure in the relationship between reasons to try and reasons to succeed. I mention them, in admittedly brief outline, to bring out the sense in which reasons to succeed are 'primary' (to use Honoré's word) as well as the sense in which they are not. Reasons to succeed are primary in the sense that the *intelligibility* of reasons to try depends on the *independent intelligibility* of reasons to succeed. If there can be no such thing as a straightforward non-derivative reason to succeed then there can be no such thing as a derivative reason to try. From this it follows that there can be no such thing as trying (or deciding, intending, aiming, etc.), and thus there can be no such thing as a non-derivative reason to try either. But once we have the possibility of derivative reasons to try on the table (by admitting that there are reasons to succeed from which those reasons derive), non-derivative reasons to try also become intelligible. Thus the case for trying is not by any means exhausted – it is not necessarily even dominated – by the case for succeeding. In that sense the story of my life as a rational agent, told as a story of endeavour, is not merely a pale shadow of the story of my life, told in terms of achievement. Both aspects have some independent significance and the full story of my life as a rational agent is the story that has room for both: it is the story of my trying (or neglecting) and my succeeding (or failing),

including sometimes (double triumph) my succeeding-by-trying and sometimes (double trouble) my failing for want of trying. The last notably discouraging case is the one that the tort of negligence latches onto: the case of someone who fails for want of trying, whose action is doubly deficient in its conformity with reasons, for not only did she not succeed; judged by the applicable measure of assiduousness, she didn't even try.

This line of argument almost fulfils Honoré's promise, I believe, and does so (I hope) in the spirit in which the promise was originally intended. What Honoré was promising was an argument according to which, firstly, our *not* having straightforward obligations to succeed would be unintelligible (so that our having them couldn't possibly be unintelligible) and according to which, secondly, our obligations to succeed would have some kind of argumentative primacy over other obligations, such as obligations to take care. The argument just ventured meets these conditions readily, save only that it softens 'obligations' to 'reasons'. So can the argument be replicated, *mutatis mutandis*, with 'reasons' hardened back up to 'obligations'? Not quite. That is because the substitution of 'obligations' makes for additional asymmetries on top of those I labelled (i), (ii) and (iii). It is perfectly possible that obligations to try might derive from non-obligatory reasons to succeed, or that non-obligatory reasons to try might derive from obligations to succeed. Thus there being obligations to try without obligations to succeed is not made unintelligible by the same knock-down manœuvres as made it unintelligible that there are reasons to try but no reasons to succeed. Obligations to try to ϕ could, after all, be obligations to act for the (itself non-obligatory) reason (*inter alia*) that one's so acting will (supposedly) contribute to one's ϕing. Still, after the foregoing considerations the ball is now firmly in the court of those who deny that we can be subject to obligations to succeed. If there are indeed reasons to succeed – and there are – is there any possible reason to doubt that they could be obligations to succeed, such that failing to do as they

would have one do would be *wrong* and hence could properly be made *tortious* in law?

An obligation is no more and no less than a categorical mandatory reason. It is categorical in the sense that it applies to people independently of their prevailing personal goals. It is mandatory in the sense that it is a reason that operates, on at least some occasions, to the partial or total exclusion of at least some countervailing reasons. Why would anybody think that these particular properties – being categorical and being mandatory – could not be possessed by reasons to succeed, now that the possibility of reasons to succeed has been established? Commonly, in my view, the following mistake tends to steer lawyers' thoughts in that direction. They think that a reason to do something counts as *mandatory* if and only if we would be justified in attaching adverse normative consequences (legally or otherwise) to its nonperformance. Since attaching adverse normative consequences to non-performance would be unfair in the case of a straightforward obligation to succeed (the thinking goes), it follows that there can be no such obligation. The problem with this line of thinking is that it wheels out the problem of unfairness much too early. Once we have established that an action is obligatory, it remains to be discussed whether the obligation should be enforced, or more generally whether people should have to bear any adverse normative consequences of its nonperformance. Possibly, as Honoré points out, further conditions have to be met before such measures would be justified, including measures to overcome the institutional fairness objection. We always admitted that there would be further internecine squabbles to come on this subject: what some would regard as sufficient to overcome the institutional fairness objection (e.g. the *Rylands* v *Fletcher* condition of an ultra-hazardous activity) others would regard as insufficient to meet that objection. But it is jumping the gun to use the institutional fairness objection as an objection to the very possibility of the obligation, when it is not yet a foregone conclusion that those

who fail to perform it will be subject to any adverse normative consequences at all. To put it simply, the institutional fairness objection is one objection, and the moral intelligibility objection is another. One cannot have two bites at the institutional fairness objection by saying that obligations to succeed are morally unintelligible because if only they existed they would have unfair normative consequences. The proper response to this alleged unfairness is to detach the normative consequences, not to deny the intelligibility of the obligation. Those who say that they cannot detach the normative consequences because they are built into the very idea of mandatoriness have simply misunderstood the idea of mandatoriness.[39] The mandatoriness of a reason lies in the fact that it operates to the exclusion of at least some countervailing reasons. Whether one is subject to adverse normative consequences in the event that one does not do as the reason would have one do is a separate – detachable – matter.

VI

To the best of my knowledge, only one serious (philosophically credible) objection has ever been raised to the proposal that reasons to succeed can be obligatory reasons. It is Kant's famous objection. According to Kant:

> A good will is not good because of what it effects or accomplishes – because of its fitness for attaining some proposed end: it is good through its willing alone – that is good in itself. ... Even if, by some special disfavour of destiny or by the niggardly endowment of stepmotherly nature, this will is entirely lacking in power to carry out its intentions; if by its utmost effort it still accomplishes nothing, and only good will is left (not admittedly as a mere wish but by the

[39] To be exact they are still in the thrall of the crudest 'sanction theories' of obligation that were decisively discredited by Peter Hacker in his famous 'Sanction Theories of Duty', in A.W.B. Simpson (ed), *Oxford Essays in Jurisprudence: Second Series* (Oxford: Clarendon Press 19xx).

straining of every means so far as they are in our control); even then it would still shine like a jewel for its own sake as something which has full value in itself. Usefulness or fruitlessness can neither add to, nor subtract from, this value.[40]

These remarks furnish the first and second premisses of Kant's curious argument against the moral intelligibility of obligations to succeed. It is an argument that dwells on the fact that obligations are *categorical* reasons (not conditional on the agent's prevailing personal goals) rather than the fact that they are mandatory reasons; thus it leaves fully open the possibility of mandatory non-categorical reasons to succeed. The argument, which is completed by Kant with the addition of a third premiss several pages later,[41] goes something like this:

(1) the only source of unconditional (a.k.a. moral) value in our actions is the good will;

(2) the good will infects not the whole of what we do but only that part of it that consists in our trying to do good;

(3) performing one's obligations is of unconditional (a.k.a. moral) value;

thus (4) there can be no obligations to succeed but only obligations to try.

This objection is sometimes confused with what I earlier dismissed as the pseudo-objection to 'moral luck'. But in fact Kant's argument is much more carefully targeted and hence withstands much more critical attention. Unlike the 'moral luck' pseudo-objection, for instance, Kant's objection has no quarrel with the objective standards of trying that are at work in the legal

[40] *Groundwork of the Metaphysic of Morals* (trans. Paton New York: Harper and Row 1964), 62.
[41] Ibid. at 68.

criterion of negligence. Indeed in premiss (1) such an objective standard is explicitly set up, namely the standard of the good will, the standard of perfect moral virtue. Not everyone, Kant agrees, is capable of meeting this standard. Some people, indeed, are moral degenerates who would not recognize a moral consideration if it slapped them in the face.[42] Relative to the baseline of their attenuated moral capacities it is a stroke of bad luck that they have the moral obligations that they have, and which they are doomed by their own degeneracy to violate. So Kant's is not a version of the 'moral luck' pseudo-objection. Rather it is an objection carefully targeted against obligations to succeed. And even regarding obligations to succeed its impact turns out to be highly selective. On closer inspection Kant's third premiss calls for a modification even by the lights of his own views. In the end he only stands by the more limited claim that it must be *possible* for performing one's obligation to be an act of unconditional value. Thus possession of a good will must be sufficient, even if not always necessary, for performance of one's duties.[43] This rules out straightforward obligations to succeed, and what I called 'tall-order' hybrid obligations (obligations to succeed by trying), but it does not rule out the short-order hybrid obligation (the obligation not to fail for want of trying), nonperformance of which constitutes the tort of negligence at

[42] Thus it is an error to associate Kant with the thesis that' ought' implies 'can' in the way in which this thesis is normally read. In the sense in which it is normally read it is taken to mean that those who lack certain moral capacities lack the corresponding obligations. Kant believed that the doctrine worked in the opposite way. He believed that since (necessarily) everyone has the obligations it follows that fundamentally they have the moral capacities as well. *Qua* human they have it in them to be less incapable than they are. As Kant spells it out: 'Ethical duties must not be determined in accordance with the capacity to fulfil the law that is ascribed to human beings; on the contrary their moral capacity must be estimated by the [moral] law, which commands categorically.' (Kant, *The Metaphysic of Morals* (trans Gregor, Cambridge: Cambridge University Press 1996), 164.)

[43] See the *Groundwork*, above note 40, at 65-66.

common law. For as we saw trying (to the legally specified degree) is sufficient, but not necessary, to perform that obligation. Thus, thanks to an important Kantian concession on the third premiss, the tort of negligence is morally intelligible to Kant. On the other hand, a strict liability tort of the *Rylands* v *Fletcher* type still is not.

I do not propose to tackle Kant's argument here. Showing that it collapses – which it does, spectacularly – is a task for another paper. Here I merely leave the argument on the file, for it seems to me that by developing Honoré's sketchy thoughts, we have now done the most important work of turning the moral intelligibility objection to strict liability back against its supporters. To achieve this what we needed to do was to bear out Honoré's claim that a story of our lives as rational agents that refuses to admit the rational salience of our successes (and failures) as well as our tryings (and neglects) is an unintelligible story. It cannot even accommodate the rational salience of our tryings (and neglects) and so, in the end, comes down to no story at all. You may say, reprising an earlier worry that we postponed, that it is one thing to establish the *rational* salience of success and another to establish its *moral* salience. Maybe there are reasons to succeed but surely they need not be *moral* reasons? Possibly in some senses of the often-unhelpful word 'moral' this is true. For example, in Kant's rather technical sense, only non-derivative reasons to try could possibly be moral reasons, whereas derivative reasons to try and the reasons to succeed from which they are derived would instead be set aside as 'prudential' reasons. Such matters of classification need not detain us here. For the notion of the morally intelligible that was implicated in the moral intelligibility objection was not, you will recall, a narrow one in which moral considerations could be contrasted with prudential ones or aesthetic ones or economic ones, etc. The notion of the morally intelligible that we had in mind was merely the notion of what is intelligible apart from the law (or similar institutional arrangements) so that the law could intelligibly claim to be

binding people from something other than a narrowly legal point of view. Our argument showed that reasons to succeed are indeed intelligible apart from the law – i.e. morally intelligible in the relevant sense – and that reasons to try are not morally intelligible, in the relevant sense, without them. From here it is, in my view, but a relatively short step to the conclusion that at least the first half of this conjunction still holds true when the reasons in question are obligatory. There is no special problem, in my view, with mandatory categorical reasons to succeed. But instead of making that short step here let me just leave it the hands of others, including those who remain attached to Kant's curious argument, to try and block it. That, I think, is the way best to honour the spirit of Honoré's ground-breaking contrarian contribution to the philosophical study of the law of torts, and indeed the philosophical study of the human condition.