

CAUSATION AND RESPONSIBILITY
RICHARD HOLTON

The notions of causation and responsibility are deeply entwined in the law. Two questions about responsibility are central to any legal system: (i) for which *consequences* of your actions are you responsible? And (ii) for which of your *actions themselves* are you responsible? Our main focus will be on the former, but causation plays a fundamental role in both.

What is the responsibility in question? It is not moral responsibility but legal. Just as legal obligation is distinct from moral obligation, so legal responsibility is distinct from moral responsibility: one can have moral responsibilities that are not legal, and, barring a general moral responsibility to obey the law, legal responsibilities that are not moral. Equally though, legal responsibility should not be equated with legal culpability; I may have a legal responsibility in an area in which I have no culpability: perhaps I have done nothing even *prima facie* illegal, or I have but I have a justification or an excuse. We can think of legal responsibility as something like answerability: I am legally responsible for those things I must legally answer for (see Duff 2007 and Gardner 2007, 2012 for different ways of developing this approach.)

Understood this way, the importance of causation in determining which consequences of my actions I am legally responsible for is obvious. The very notion of consequence is itself a causal one: by and large, the things for which one is responsible, whether in criminal law or in torts, are those that causally issue from one's actions. Sometimes the causal element is explicitly noted in the law; sometimes it is implicit in the causal terms that the law uses. Notions such as 'murder', 'battery' or 'negligence' all involve causal relations: a necessary condition for murder is that the defendant *caused* the victim to die.

So should we understand causation as a *necessary* condition on responsibility for outcome? That rather depends on how we understand causation. Take, for instance, a case of overdetermination: you and I each do something that would have been enough on its own to bring about the effect. It might seem that each of us should be responsible for it, but on some prevalent accounts of causation—most notably those which say that a cause is something without which the effect would not have occurred—neither of us looks to count as a cause. So either we have to give up on the idea that causation is necessary for responsibility, or we have to give an alternative account of causation. A great deal of the discussion here has been devoted to trying to formulate an alternative account that will do the work.

A second issue concerns whether having caused an outcome is *sufficient* to be legally responsible for it. In an unqualified form that claim is surely too strong. Other people, or rogue happenings, might intervene in the chain of events, and that might affect responsibility. Moreover, if causation is, as many think, a transitive notion—

so that if A causes B, and B causes C, it follows that A causes C—then the consequences of one's actions extend indefinitely into the future. Are we responsible for all of them? Or only some, those that are sufficiently close to the action, those that have the requisite proximity? And if we take the latter course, how do we characterize 'sufficiently close'? Is that itself to be characterized in terms of the metaphysics of causation—in terms of 'cause-in-fact' as it is often called? Or is it rather a normative question, to be settled by reference to other factors about what we ought to be held responsible for? Much of the debate around the legal notion of causation has focused on how to draw the boundary here.

Turning to the role of causation in determining which actions we are responsible for, we find that parallel issues arise. Most jurisdictions recognize some kind of insanity defence. If someone kills while suffering from paranoid schizophrenia, they may not be legally responsible for what they have done. But in most jurisdictions, it is not enough that they kill and that they suffer from the illness: it needs to be shown that the illness was relevantly implicated in the killing. The natural way to understand this is again in terms of causation: there has to be a causal influence of the illness on the killing. If so, many of the issues that arose for in determining responsibility for outcome—issues around preemption, overdetermination and proximity—will have parallels here.

We will look at these two issues in turn, with our focus on responsibility for outcome, but first it will be helpful to examine ways of understanding the metaphysical framework surrounding the notion of causation.

THE METAPHYSICAL FRAMEWORK

The philosophical literature on causation is enormous and increasingly complicated; there is no chance of summarizing it. Still, two debates can be identified which will be enough to structure the discussion.

The first debate is over which kind of relation captures the essence of causation. There are two distinct ideas that underpin much of our everyday thinking. One is that a cause is *sufficient* for its effect: the cause is whatever it takes to bring the effect about. The other is that a cause is *necessary* for its effect: that in the absence of the cause, the effect would not have happened. Philosophical *analyses* of causation—that is, accounts that seek to define causation in terms that do not mention causation themselves—typically take one or the other as fundamental. For many years, particularly under the influence of David Hume, philosophical theories focussed on the sufficiency side. But since the 1970s the philosophical orthodoxy has shifted radically, so that now most accounts focus on necessity: causes are those things without which the effect would not have happened.

Part of the explanation for that shift is tied up with the second major philosophical debate around causation, which concerns the metaphysical machinery to be used in

the account. In characterising causation, should we make use of modal notions, most centrally of counterfactuals, statements that make reference to how things *would have* turned out under other circumstances? Or is such talk philosophically illegitimate? Here again Hume's work casts a long shadow. Hume was suspicious of talk that didn't concern how things can be observed to be; modal notions tended to fall prey to this suspicion. So his version of the sufficiency claim was formulated in terms of regularity: very roughly, C causes E if and only if things of the same type as C are always followed by things of the same type as E. Such an account makes no mention of how things would have been (although oddly Hume does at one point say that it is equivalent to a counterfactual formulation). After Hume, many have realized that such an account is vulnerable to various kinds of counterexample (What if C is the only one of its type? What if C and E are both effects of some common cause? what if C has been invariably followed by E as a result of some great coincidence?) and so later empiricist accounts, from Mill to Mackie, typically made use of the idea of a natural law—C causes E iff there is some *law-like regularity* linking C and E—where the idea of a law is not itself defined in counterfactual terms.

In contrast to all this, accounts that focus on the necessity of the cause for the effect have typically been formulated using explicitly counterfactual machinery. The most influential are from David Lewis, whose accounts of causation are couched in terms of the semantics of counterfactuals that he developed: again roughly, a counterfactual of the form 'If it were the case that P it would be the case that Q' is true just in case the closest possible world (i.e. closest to the actual world) in which P obtains is a world in which Q obtains. (Closeness is to be understood in terms of similarity.) Using this account, Lewis defined a notion of *causal dependence*: E is causally dependent on C if and only if (i) if C were to occur, then E would occur, and (ii) if C were to fail to occur, then E would also fail to occur. Causation itself was then analysed by Lewis in terms of a chain of causal dependencies (Lewis 1973). C causes E just in case either E causally depends directly on C, or there is a chain of dependencies between events more or less long stretching back from E to C. This guarantees that causation will be a transitive notion. (Lewis came to refine his account in significant ways, but this won't concern us here; see Lewis 2000.)

However, this association of sufficiency with regularity accounts, and of necessity with counterfactual accounts, is largely an historical accident. Lewis, as we have just seen, defines causal dependence in terms of two clauses: whilst the second (If C were to fail to occur, E would fail to occur) is clearly a necessity condition, the first (If C were to occur, E would occur) looks like a sufficiency condition. In fact, on Lewis's account of counterfactuals, that condition is trivially satisfied when C and E have both happened (the closest world to the actual world in which C obtains is the actual world itself, and that is a world in which E obtains), but on other accounts of counterfactuals—those that reject 'strong centering'—it will not be. On such accounts C is sufficient for E just in case in *every close world* to the actual world in which C happens, E happens—here Hume's quantification across different actual events is replaced by quantification across different occurrences of the same event in different possible circumstances. Or, more plausibly, since most causes will require

other factors to be present if they are to be sufficient we might say that something is a cause just in case it is a non-redundant member of a set of conditions that suffice for the effect; we will return to such an account—propounded by Mackie, and in the legal literature by Wright—shortly.

We will not be further concerned with regularity accounts here, or of the empiricist scruples that motivate them. Attempts to dispense with all modal talk and to replace it with notions of law-like regularity or of entailment, where these are not in turn cashed out in modal talk, have not been very successful (Paul & Hall 2013 pp. 42–3). Moreover empiricist scruples themselves are far from natural. Ordinary speech is full of counterfactual constructions, as is most legal practice and legal theorizing. So the focus here will rather be on whether necessity or sufficiency conditions should be central—or whether we need some mixture of the two—where both of these are to be understood in explicitly modal counterfactual terms. Before examining how this has played out in the law, let us first examine the problems that the two approaches face.

PROBLEMS FOR THE ACCOUNTS

Problems for necessity accounts

The first problem facing necessity accounts is that they *undercount* causes, because of how they handle cases of redundancy: cases in which there are other causes that would bring the effect about if the cause in question didn't do so. Such cases are typically divided into two types:

- (i) cases of *preemption*, where it is intuitively clear that there is one factor that is the cause—the preempting cause—but where something else would have caused the same (or a very similar) effect if the preempting cause had not already done so: assassin A shoots V before assassin B can do so. These are further divided into cases of *early* pre-emption, where the preempting cause stops or interrupts the alternative cause—hearing A's shot, B does not fire; and cases of *late* pre-emption, where the alternative cause is not interrupted but simply happens too late—A's shot kills V before B's shot hits.
- (ii) cases of *overdetermination*, where there is redundancy and yet there is a symmetry, so that each putative cause does intuitively seem to contribute to the effect. These are further divided into cases of *sufficient overdetermination* where each of the causes would have been sufficient on its own, and so each of which is redundant (assassins A and B both shoot V simultaneously, each is enough to kill), or cases of *insufficient overdetermination* where the causes, whilst being individually redundant, are not individually sufficient, since several but not all of them are required for the effect (assassins A, B and C all shoot V simultaneously; two shots are needed to kill).

Necessity accounts clearly do badly on all of these redundancy cases: in their simplest form they will conclude that there is no cause, because in each case no single factor is necessary.

A second problem with necessity accounts concerns the *over-counting* of causes. Everything that is necessary for an effect counts as one of its causes: not just the action of the arsonist in striking the match, but also the presence of oxygen, the action of her parents in conceiving her, and the origins of the universe. Perhaps more counterintuitive still, *absences* seem to count as causes, and, again, these may be very distant: so not just the failure of the police officer to check what the arsonist was doing, and the absence of a sprinkler system, but also the failure of a thief to steal her car thereby stopping her getting to the site, and, more distant still, the failure of other suitors to have stopped her parents ever conceiving her. Philosophers tend to be rather dismissive of such worries, happy to say that, once we stop talking about *the cause* of an event and start talking about *causes*, all can be welcomed in. But in the legal world, where cause brings responsibility, such liberality is a problem.

Problems for sufficiency accounts

Sufficiency accounts also face substantial problems, some of their own, and some that they share with necessity accounts. The first concerns how to understand the relevant notion of sufficiency. Obviously what we think of as a cause is not sufficient on its own: pulling the trigger would not have caused the gun to fire if there had not been oxygen present to react with the explosives in the cartridge. So a first move, as with necessity accounts, is to think in terms, not of the single cause, but of a set of causes. We might try to formulate a sufficiency account by saying that something is a cause if it is one of a set of factors that are sufficient for the effect. But that is too liberal: each of the factors in the set had better do some work. If we add this requirement that they do some work, saying that something is a cause if it is a *necessary* member of a set of factors that are jointly sufficient for the cause, we arrive at the kind of account that was introduced into legal theory by Hart and Honore and has been most influentially defended in the philosophical literature by Mackie (under the heading INUS: Insufficient but Necessary part of a condition which is itself Unnecessary but Sufficient). This of course builds in elements of necessity as well as sufficiency. Moved by empiricist scruples, Mackie tried to account for both in terms of laws; but both can more naturally be analysed in terms of counterfactuals.

How does a mixed sufficient account along these lines fare with the redundancy problems that beset the necessity account? In cases of preemption, where A shoots and kills V before B has the chance to do so, it has the mirror image problem. The sufficiency account has no difficulty classing A's shooting as a cause of V's death: A is a member of a set of factors that are sufficient for it. Its problem comes with the preempted event: since B's shooting is also a member of a set of factors that are sufficient for V's death, it will wrongly identify it as a cause. (There is some debate about this—some (Strevens 2007; Moore 2009) have argued that it will not count, since V will not be alive to be killed; but for an effective response see Maslen

2012.) When it comes to overdetermination, sufficiency accounts do better. In cases of sufficient overdetermination (A and B each shoot V simultaneously; each shot would be enough for V's death), both A's shooting and B's count as causes, since each is a necessary part of a sufficient set: A's of a set that does not include B's, and B's of a set that does not include A's. They likewise do well with cases of insufficient overdetermination (A, B and C all shoot, any two shots would be enough for V's death): each counts as a cause since, for instance, A's shooting is a necessary member of a sufficient set that contains just A's shooting and B's shooting, and so on for the others.

With respect to the issue of counting everything in the causal history of an event as a cause, INUS-style sufficiency accounts are in much the same boat as necessity accounts: since any set that contains the sufficient causes must contain *all* of the conditions needed for the outcome, very distant conditions will count. Indeed, given that necessity accounts typically deny causation in cases of overdetermination, sufficiency accounts will tend to include more things under their tally. (Whether it will contain all of the relevant omissions and absences depends on quite how sufficiency is understood: if it is phrased in terms of counterfactuals, they will be included; but if it is phrased in terms of natural laws, the question will hinge on whether the relevant laws will need to mention them.)

The philosophical literature looking at ways of avoiding these problems for the two accounts is enormous. It is safe to say that there is no agreed solution. (For a recent review of options see Paul and Hall, 2014.) As with so many philosophical issues, the thought arises that there may not be an analysis of the notion of causation. Perhaps the notion is primitive; after all, on pain of regress, some notions have to be. If so, the two accounts might be better viewed as heuristics than analyses; we can look to that in the context of the legal practice.

ACCOUNTS OF CAUSATION IN THE LAW

Far and away the most prevalent account of causation in the law, regularly invoked both in courtroom practice and in legal writings, is a necessity approach: the *but-for* or *sine qua non* test. But in the light of the problems already mentioned, its inadequacy given a simple formulation has been widely recognized: "It is now, I think, generally accepted that the "but for" test does not provide a comprehensive or exclusive test of causation in the law of tort" said Lord Bingham in *Chester v. Afshar* [2005], going on to cite both issues of undercounting and over-counting.

So what should stand in its place? Much academic writing, though less courtroom practice, has focussed on sufficiency accounts, especially as given by Richard Wright in his NESS (Necessary Element of a Sufficient Set) formulation, which follows Mackie's INUS account fairly closely (for discussion of the differences, and responses to criticisms, see Wright 2011). Again though, as we have seen, sufficiency accounts

have problems with redundancy, so, despite the amount of work devoted to their development, it is far from clear that they can work on their own.

One response is to produce a new account by disjoining the two rivals: something counts as a cause if it passes *either* the *but-for* test, *or* the NESS test (see Steel, 2015 for a proposal along these lines). This helps where one test is accurate but the other *under-counts*, as with overdetermination. But it doesn't help if one of them *over-counts*, as happens with preemption: since the necessity test is positive for both the preempting cause and the preempted potential cause, the disjunctive account will classify both of them as causes.

A more nuanced approach makes use of the tests selectively: using one or the other, suitably modified if need be, in the cases where we have reason to believe that they will work well. Compare our use of common-sense tests in other domains. We may test for the presence of water on the path by checking if it feels wet. But if the temperature is below freezing that test will not work; now we check if it feels slippery. Indeed, if it feels wet at such a temperature, then that's good grounds for thinking that any liquid present *isn't* water. Conversely, if we are not sure what the temperature is, but are confident that the liquid on the ground is water, we might use its feel to estimate the temperature. And if we know that there are interfering factors—perhaps someone has spread salt on the path—things change again. We employ tests defeasibly and flexibly, amending our beliefs on the outcome of the tests we deploy, but also amending the tests we deploy on the basis of our beliefs.

Likewise, it is plausible that the necessity and sufficiency criteria should be seen as *defeasible* tests for causation. In a case of overdetermination, the sufficiency test will work better than the necessity test. In a case of preemption, neither test will work well, so we will need to look to other tests; but the fact that the necessity and sufficiency tests give such divergent results is itself an indication that preemption is involved (to get an idea the kinds of further tests that may help here, see Paul and Hall 2014—though there they are formulated as candidate analyses). On such an approach we cannot *define* cause in terms of one of these tests: they could not be used to introduce the notion of causation to someone who had no understanding of it. It is our prior grip on the notion of causation that enables us to see where the tests can be usefully employed. But it doesn't follow that the tests have no function. They may work, as we have seen, as heuristics, enabling us to identify cases of causation in many concrete cases: this is plausibly how a jury might deploy them. Moreover, they might help us articulate the reasons for coming to conclusions in particular cases, and then to build up a body of precedent to deal with future cases. (For some an approach along these lines see Thomson 2007; for the general issue of whether rules in law work as decision procedures that could be used by someone who had no prior grip on the issues, or whether they serve to articulate our implicit reasoning and make them useful for future application, see Dewey 1924)

What though of the problem of counting everything in the causal history of an event as a cause in equal standing, no matter how distant from the event, or how insignificant its contribution appears to be? We focus first on the latter of these two

issues, the need to estimate the degree of contribution that a given factor makes to an outcome.

The legal need for such a test has been somewhat obscured by two features of legal practice. First, in many jurisdictions, the degree of causal contribution is irrelevant to criminal responsibility: even someone who merely aids or abets the commission of a crime is as responsible as the principal offender. Second, within torts, the widely applied doctrine of joint and several liability means that each of the defendants will be liable up to the full amount of the obligation: the plaintiff may recover damages from any one of the defendants, leaving it up to that defendant to recover the due share from the others. These doctrines have had the consequence that it has been somewhat less important to determine the degree of causal contribution than might be expected. But causal contribution is not legally irrelevant. For a start, not all jurisdictions embrace joint and several liability; and even those that do need some way of address the apportionment of liability amongst the defendants if one of them goes on to sue the others.

We concentrate on cases of insufficient overdetermination. Such cases are real in tort and criminal law: pollution by many manufacturers causes illness or damage, but no one polluter's contribution was necessary; the members of a company board vote unanimously not to recall a dangerous product, where a simple majority would have been sufficient to carry the policy (see Stapleton 2013 and Steel 2015 for a discussion of such cases). Here courts tend to talk of 'contribution' rather than cause, a graded notion rather than all-or-nothing; but how is this to be understood?

In cases where direct physical forces are involved, we can measure the contribution in terms of those forces. If A, B, and C push a car off a cliff, or hold another person down, we can measure their physical contributions, even if the total force is greater than is needed (Stapleton 2013). Perhaps the same reasoning can be applied to polluters: we can measure the amount of pollutant introduced by each of the defendants. But it is hard to see how to do this for the votes of the board: there are no physical forces involved there. If the board members all have the same voting power it is natural to think of their contributions as equal, but this does not have to be the case: some could be assigned more votes than others. Other cases will be even harder to assess: a suicide may be over-determined by the bullying and harassment that led to it, yet the contributions of those involved may differ.

Two approaches have been suggested. One involves counterfactuals: an attempt to measure how big a contribution is made by each factor to the change from the world in which the effect does not take place, to the world in which it does (for different approaches see Chockler and Halpern 2004; Braham and van Hees 2009). The other involves probabilities: how great an impact does each factor have on the likelihood that the effect will take place (Kaiserman 2017). Again though the chances of providing a reductive analysis of the notion look bleak. Counterfactual accounts require a prior individuation of the relevant events if they are to get a grip: they are better suited to measure contributions in the case of the voters than of the suicide. Probabilistic accounts avoid this problem, but face others. In general a

factor can fortuitously cause an outcome even though it lowered the chance that that outcome would happen (Hitchcock 1995); likewise there will be cases where a factor can be non-redundantly involved causing an outcome even though it lowered the chance of that outcome. Again though, this does not show that these accounts have no place. It suggests rather than they should be used as *defeasible tests* for contribution to causation, rather than as reductive analyses of it.

Such accounts might also help shed some light on the other problem that faces both the necessity and the sufficiency accounts: where to draw the line between causes that are sufficiently close to the effect that the author of the cause bears responsibility for it, and those that are not. This issue, that of proximity or remoteness, the idea that in some sense remoteness decreases liability, remains the subject of intense legal controversy.

Clearly we should not take the idea of spatial or temporal proximity too literally. A defendant who acts, using the appropriate technology, at great distance, does not thereby avoid liability—a defendant who places a bomb on a plane so that it explodes on the other side of the globe, or fits it with a slow fuse so that it explodes after many months, is just as liable as one who detonates it here and now. Nor does the complexity of the causal process in itself make a difference: using a sophisticated detonator that involves many causal stages is no different to directly lighting a simple fuse. The distance should rather be causal distance.

One approach applies the idea of degree of causation directly here: if causes come in degree, then we might think that responsibility only attaches if an agent's causal contribution is sufficiently large. We might think of cause as diminishing to zero over enough steps; this would be to deny its transitivity (Moore 2009). Or we could maintain the idea that cause is a transitive notion, while saying that having a large causal influence is not; what we need to do is quantify the size needed for responsibility.

The alternative (probably more widespread) approach takes the view that proximate causation is concerned with other issues, to be contrasted with the kinds of metaphysical issues that we have looked at so far. That is, it is not a species of 'cause in fact', but a specifically legal notion (Stapleton 2008). We need to distinguish two different ways of spelling this out. The first is that proximate causation is badly named: it has nothing specially to do with causation, as ordinarily understood, at all. Rather it is to do with some other issue, most plausibly something to do with the intention or foresight of the defendant. The second view is that proximate causation is indeed centrally concerned with causation, but that it involves a normative constraint on which causal chains bring responsibility. These constraints might themselves need to make use of causal distinctions; but they are not distinctions that would naturally stem from metaphysical considerations independently of their normative significance.

On the first approach, responsibility runs out if the agent did not intend, or could not have reasonably foreseen, the consequences of their act (the details here would

depend on the *mens rea* attaching to the particular offence). Or, if the outcome was foreseen, whether the way it came about was foreseen: whether there was the appropriate 'harm within the risk'. An attempted murder would not be a successful murder if the death of the victim resulted, quite unexpectedly, from their suicide in response to the attempt. (Moore 2009) However, once proximity is understood in this way, the question arises whether it is needed as a separate category: why not let the *mens rea* requirement do all of the work?

On the second approach proximity is understood in terms of the specific kinds of causal chain that are needed for legal responsibility. A central consideration here is the one that Hart and Honoré drew from their discussion of a large body of case law: that responsibility is blocked if other agents, or some freak events, supersede the initial cause. Take one of their examples: suppose D, aiming to kill V, knocks him to the ground, where he is killed by the freak falling of a tree, or by the actions of some third party, who, seeing V helpless on the ground, maliciously fells the tree onto him. D would not be responsible for V's death.

Hart and Honoré take this to show that in such circumstances D would not have caused V's death, but this is implausible. To see this, note that in both cases, lack of intention on D's part is still required if D is to avoid responsibility. If D's plan had involved the falling tree, or the expectation that the third party would become involved, it would still be a case of murder. But how can the issue of whether or not D was the cause of V's death turn on issues of D's beliefs? (Thomson 2008; for a broader criticism to the same conclusion see Wright 2008) Much more plausible is the thought that causation is present in all cases, but that only certain types of causation are sufficient for responsibility.

Similar ideas can be applied to the issue of complicity. In cases where being an accomplice rather than a principal is important, the difference is not to be understood in terms of the principal being the cause of the offence whereas the accomplice is not. If a would-be accomplice has no causal impact they are not an accomplice (although their impact might be over-determined, and so badly assessed using necessity tests). It is rather that the accomplice causes the principal to cause the offence, and not vice versa. (See Gardner 2007, Ch. 3)

The point is equally plausible in the case of omissions. As we saw, both necessity and sufficiency accounts find it hard to distinguish between the causal strength of acts and that of omissions; both tend to count equally as causes, and it is hard to see how that could be avoided (though see Moore 2009 for a valiant attempt). Nevertheless, in general Anglo-American law is reluctant to prosecute for omissions except those that contravene an established duty to act. (For general discussion here, see Alexander 2002.) While the philosophical attempt to delineate general principles that distinguish acts from omissions (or distinguish doing from allowing, which is clearly not the same thing) has certainly not reached consensus (see Steinbock and Norcross 1994 for a sense of the difficulties) most approaches look to establish different sorts of causal chains, rather than to class omissions as fundamentally not causal. Recent attempts to distinguishing those omissions that

are relevant from those that are not have made use of contrastive accounts: accounts that make reference to salient alternatives amongst those that could have happened if the cause and effect had not taken place. (Schaffer 2010)

WHICH ACTIONS ARE WE RESPONSIBLE FOR?

So far we have examined the role of causation in the attribution of responsibility *downstream* of the agent's action; in closing we briefly examine its role *upstream* of it.

To fix ideas, consider the case of the right-wing terrorist Anders Breivik. In 2011 he detonated a bomb in Oslo that killed eight, and then shot dead 69 young supporters of the Workers' Youth League at an island summer camp. An initial psychiatric assessment found him insane on the grounds of paranoid schizophrenia. Under Norwegian law that would have been enough to provide a defence, even though no evidence was given that the insanity was the cause of his actions—it was not argued, for instance, that he thought the children were trying to kill him. (A second psychiatric assessment came to a different conclusion, judging him as suffering from personality disorders, hence not psychotic, and so not insane under Norwegian law; the court agreed.) On the Norwegian approach, insanity defences are a form of *status* excuse, like that conferred in Anglo-American law on children under the age of 7, who are held incapable of being criminally responsible. On such an approach, the insane are simply not taken to be moral agents in the first place; issues of causation do not enter in.

Such an approach is very foreign to most Anglo-American accounts of insanity (though see Moore, 2015 for a contrary view). Instead the standard approach is to think that the illness must be causally involved in the act if it is to remove responsibility. (The same is typically true for justifications and excuses: thus it is not enough for a partial defence of provocation that the accused kills having been provoked; it must be that the provocation causes them to lose self control which in turn causes the killing; Holton and Shute 2007.) Given this, we might expect to see many of the issues discussed so far arising here. What if an act is overdetermined by the illness and something else? What if the casual effects of the illness lack proximity? That such issues have not been to the fore in legal contexts probably shows the difficulty of establishing the causal effects of mental happenings. As neuro-imaging becomes more sophisticated, we may start to see it more.

BIBLIOGRAPHY

- Alexander, L. 2002: 'Criminal Liability for Omissions: An Inventory of Issues' in S. Shute and A. Simester (eds.) *Criminal Law Theory: Doctrines of the General Part* (Oxford: Oxford University Press) 121–42.
- Braham, M. and M. van Hees 2009: 'Degrees of causation' *Erkenntnis* 71, 323–44.
- Chockler, H., and J. Halpern 2004: 'Responsibility and blame: A structural-model approach' *Journal of Artificial Intelligence Research* 22, 93–115.
- Dewey, J. 1924: 'Logical Method and Law' *The Philosophical Review* 33, 560–72.
- Duff, A. 2007: *Answering for Crime: Responsibility and Liability in the Criminal Law* (Oxford: Hart Publishing).
- Gardner, J. 2007: *Offences and Defences* (Oxford: Oxford University Press)
- 2011: 'Relations of Responsibility' in R. Cruft, M. Kramer and M. Reiff (eds.), *Crime, Punishment and Responsibility* (Oxford: Oxford University Press) 87–102.
- Hart, H.L.A. and T. Honoré 1985: *Causation in the Law*, Second Edition (Oxford: Clarendon Press).
- Hitchcock, C. 1995: 'The Mishap at Reichenbach Fall: Singular vs. General Causation' *Philosophical Studies* 78, 257–91.
- Holton, R and S. Shute 2007: 'Self-Control in the Modern Provocation Defence' *Oxford Journal of Legal Studies* 27, 49–73.
- Kaiserman, A. 2017: 'Partial Liability' *Legal Theory* 23, 1–26.
- Lewis, D. 1973: 'Causation', *Journal of Philosophy* 70, 556–67.
- 2000: 'Causation as Influence', *Journal of Philosophy* 97, 182–97.
- Mackie, J. L. 1965: 'Causes and Conditions', *American Philosophical Quarterly*, 2 245–64.
- Maslen, C. 2012: 'Regularity Accounts of Causation and the Problem of Pre-emption: Dark Prospects Indeed' *Erkenntnis* 77, 419–34.
- Moore, M. 2009: *Causation and Responsibility* (Oxford: Oxford University Press)
- 2015: 'The Quest for a Responsible Responsibility Test: Norwegian Insanity Law After Breivik', *Criminal Law and Philosophy*, 9, 64–93.
- Paul, L. A. and N. Hall 2013: *Causation, A User's Guide* (Oxford: Oxford University Press).
- Schaffer, J. 2010: 'Contrastive Causation in the Law' *Legal Theory* 16, 259–97.
- Stapleton, J. 2008: 'Choosing What We Mean by "Causation" in the Law', *Missouri Law Review* 73, 433–80.
- 2013: 'Unnecessary Causes' *Law Quarterly Review* 129, 39–65.
- Steel, S. 2015: *Proof of Causation in Tort Law* (Cambridge: Cambridge University Press).
- Steinbock, B, and A. Norcross (eds.) 1994: *Killing and Letting Die*, 2nd edition, New York: Fordham University Press.
- Strevens, M. 2007: 'Mackie Remixed', in J. Campbell, M. O'Rourke, and H. Silverstein (eds.) *Causation and Explanation* (Cambridge MA: MIT Press, 2007).
- Thomson, J. J., 2007 'Some Reflections on Hart and Honoré, Causation in the Law' in M. Kramer, C. Grant, B. Colburn, and A. Hatzistavrou (eds.) *The Legacy of H.L.A. Hart: Legal, Political, and Moral Philosophy* (Oxford: Oxford University Press) 143–64.

- Wright, R. 2001: 'Once More into the Bramble Bush: Duty, Causal Contribution, and the Extent of Legal Responsibility', *Vanderbilt Law Review* 54, 1071–1132.
- 2007: 'The Nightmare and the Noble Dream' in M. Kramer, C. Grant, B. Colburn, and A. Hatzistavrou (eds.) *The Legacy of H.L.A. Hart: Legal, Political, and Moral Philosophy* (Oxford: Oxford University Press) 165–77.
- 2011: 'The NESS Account of Natural Causation: A Response to Criticisms' in R. Goldberg (ed.) *Perspectives on Causation*, (Oxford: Hart Publishing) 285–322.