NO SUCH THING AS ACCIDENT: RETHINKING THE RELATION BETWEEN CAUSAL AND MORAL RESPONSIBILITY

Mark R Reiff
University of California at Davis
(mreiff57@gmail.com)

Canadian Journal of Law and Jurisprudence 28 (2015): 371-397

In my first year torts class, when I get to the section of the course that deals with negligence, I often begin with the following hypothetical. Suppose you are sitting in your office and you get a call from an old friend from law school who has a potential litigation matter that he wants to refer to you (he does not handle litigation matters himself). He has a client whose gainfully employed 25-year old son walked into a newsagent to buy a candy bar, and something happened—he apologizes that he does not know exactly what—but as a result of this the son ended up a quadriplegic. Your friend knows no other details, and what's more, he says he is about to go on vacation, and therefore has to get the matter of representation settled immediately, so you have to tell him right there and then whether you would be willing to take on the case or your friend will have to refer the matter to someone else. But don't worry—he has already spoken to the father about your fee, and has been assured that the son is prepared to sign a standard contingency fee agreement, under which you will be entitled to keep one third of any recovery you obtain on his behalf. What do you tell him?

At this point, I usually put the matter to a vote, and the results are pretty consistent from year to year. There is always a smattering of students, usually two or three but never more than five, who say they would take the case. But everyone

else—and this means over three hundred students—would refuse. And the reasons they give for why they would refuse are exactly what you would expect. They need to know much more about what happened in the store, they say, in order to know whether anyone was at fault, and if so, who? Without more details, they simply cannot make a judgment on whether there is any legal liability on behalf of the store owner (or anyone else for that matter), and so they are not in a position to estimate if there is real chance of recovering damages here, and if so, how much. It would be both irresponsible and financially foolish to take this case on a contingency basis without such information, they say, so the only thing they can do is refuse.

Of course, there is much in their textbooks that supports my students' view. After all, in the section on negligence law, a great deal is made of the element of fault. There is no claim for negligence without fault, which is the kind of claim you would want to bring here, and textbooks are filled with cases about whether the relevant standard of care has been met.[1] As a result, students understandably get the impression that the outcome of many cases turns on this. So how could I expect them to commit to this client without knowing more about what went on, for only then could they form an opinion as to whether the relevant standard of care was breached and who may have breached it. After all, they say, *it could have simply been an accident*, for which no one is at fault.

[1] See, e.g., *Nettleship v Weston* [1971] 2 QB 691, a standard inclusion in English tort law casebooks and textbooks in the section on standard of care. In that case, despite driving at a slow rate of speed and being instructed exactly what to do, a learner driver managed to drive up over the curb, onto the sidewalk, and crash into a lamp post, injuring her passenger, the driving instructor. The issue on appeal was what standard of care should be applied—that of a reasonable driver or that of a reasonable learner-driver. As if this could have possibly made a difference. Even at the lower standard of care, it is difficult to believe that there was no fault here. Otherwise, we all take our lives in our hands anytime we go anywhere we may encounter learner-drivers on the road. For examples of the attention given the standard of care in American textbooks, see an B. Dobbs, *The Law of Torts* (St Paul, MN: West, 2000) §§116 et seq.

Nonsense, I tell them. This is a great case. You would be crazy not to take it, and here's why. In real life, where litigation is only cost effective when a serious injury is at issue, *fault does not matter*. At least fault does not matter very much. What matters is damages and causation. Damages matter for two reasons. First they matter because it is only cost effective to take on a case if damages are large. If the potential damages are small, it does not matter what happened, it is simply not worth bringing the elaborate mechanism of the justice system to bear on the parties involved. And if the potential damages are large, so it is worth bringing the elaborate mechanism of the legal system to bear, there is going to be no problem proving fault. For whenever someone is seriously injured, it must be *someone's* fault (and by "fault" I mean something for which someone is morally responsible). It may be the plaintiff's *own* fault, at least to some extent, but there is no such thing as an incident in which someone is seriously injured and *no one* is at fault. People simply do not want to believe you can walk into a newsagent to buy a candy bar and end up a quadriplegic without someone failing to exercise reasonable care.

This is simply a reflection of human nature. If we lived in a world where people could be seriously injured in circumstances where no one was at fault, the world would be a very scary place to live. Danger would be lurking everywhere, and there would be nothing anyone of us could do to fully protect ourselves from it. In contrast, if we live in world where as long as everyone exercises reasonable care, no one can get seriously injured, well, that is a very comforting world. That is the kind of world I would like to live in, and I dare say it is the kind of world we would all prefer to live in if we had the choice. Because people do not want to live in that very scary world where they can be seriously injured even if everyone exercised reasonable care, they are going to endow the reasonable person with abilities that most people

actually do not possess, and so it is almost never going to be the case that you cannot prove fault.[2] I'm not saying that it will always be easy—I am merely saying that you start out with such a huge advantage that if you know what you are doing you are going to be able to prove fault. Your client injured himself when he tried to reach for a box of cereal on a high shelf? Why wasn't there a sign warning customers not to try this on their own, or a small stool to stand on like they have in libraries if you need to get to an upper shelf? If there was a stool and your client slipped off, why wasn't there a slip resistant coating on the top, or why wasn't the stool more stable? If your client was struck by lightning in the store, why wasn't there a lightning rod on the building? If he was shot in the back by robbers who came into the store after him, why did the store owner simply not give the robbers what they wanted? Why did he not use video surveillance to discourage robbery attempts? And so on. Indeed, whatever factual scenario you can come up with, no matter how outlandish, I can come up with some reason why the store owner could have done something more to prevent the incident that is no more outlandish than the initial statement of what happened itself. I'm not saying that store owner can always or even often be shown to be *solely* at fault, but it is almost inconceivable that there could be some set of circumstances in which he or someone else with insurance coverage other than your client was not at least partially at fault.

Causation, in turn, matters because while people do not want to believe that they could be seriously injured even though no one was at fault, they *are* willing to believe that a particular person's acts or omission did not *cause* a particular injury. This does not make the world a scarier place. So causation is always something to be

[2] For further discussion of this phenomenon and some of the empirical studies examining it, see Jeffrey J. Rachlinski, "Misunderstanding Ability, Misallocating Responsibility," *Brooklyn Law Review* 68 (2003): 1055-1091.

concerned about in every case. But causation is itself a curiously complex topic. There are many, many causes of every serious injury, and therefore many people to whom we could assign *some* causal responsibility. But we do not attribute causal responsibility to everyone who in some sense or another is part of the causal chain that led to the injury. We are much more selective than this. And when we do decide who is causally responsible, we then supposedly narrow this group down even further by asking who within this group is also morally responsible. We do this, in turn, by asking who within this already select group was at "fault," meaning who within this group acted intentionally or recklessly or at least failed to exercise reasonable care, an inquiry I will refer to hereinafter as a search for "traditional fault." But if I am right about our fear of finding that someone who was causally responsible was not morally responsible too, what work is the concept of traditional fault really doing for us? Why is it that we see a finding of traditional fault as so essential to an attribution of moral responsibility? If a finding of fault, here meaning moral responsibility but not necessarily traditional fault, effectively follows automatically from a finding of causal responsibility, is not this finding of causal responsibility doing all the real moral work? Is not a finding of causal responsibility really a finding of moral responsibility too? And if it is, how can we possibly explain this?

Now some of you who think of yourselves as normative philosophers are no doubt already rolling your eyes and thinking that these are just the crude questions of a lawyer that have may some practical significance but are of no deeper philosophical interest or importance, for that is indeed a not uncommon initial reaction among those who have attended presentations of this paper. But the examination I intend to undertake is not simply an exercise in exploring the nature and consequences of the practice of litigation and the assignment of legal responsibility, something that might

be of interest to those who are engaged in such an activity but not to those who see their role as contemplating and clarifying the role of morality in the regulation of our common social life. On the contrary, what the example with which I began this paper shows, I believe, is that we *as philosophers* have the nature of morality wrong. We think that traditional fault matters much more to an assignment of *moral responsibility* than it actually does. In actuality, I contend, it is causation, not the determination of traditional fault, that lies at the heart of our moral and not just our legal judgments and practice.[3] And that is what this paper aims to prove.

I. SOME PRELIMINARY CLARIFICATIONS REGARDING THE NATURE OF MY ARGUMENT

Before I set out to do this, however, let me make clear the limiting principle under which I will be proceeding. In this paper, I will be talking only about claims of serious personal injury. Different rules may apply to claims of less serious personal injury or to injury to property or economic injury alone.[4] I also do not mean to include cases of serious injury or death that come at the natural (in terms of years) end of life. First, because damages in these cases are almost always going to be slight and such cases would therefore be excepted from the reach of my argument in any event; and second, because even in those unusual cases where for some reason damages at

[3] Note that I take no position in this paper on the metaethical question of whether there are such things as moral facts, for nothing in my argument turns on this. True, the phrasing of my argument might need to be slightly different if there are such facts, but the substance of my claims would not change. If there are moral facts, then what I say about moral and causal judgments should be understood to include such judgments and the moral and non-moral facts on which they are based; if there are no moral facts, then what I say about moral and causal judgments should be understood to include such judgments and the non-moral facts on which they are based, which in that case would be the only kind of facts there are. When I refer to our moral or legal practice, in turn, I am referring to how we arrive at moral and causal judgments within that particular practice, whichever that may be.

[4] Given this limitation, most of the cases that I will be discussing will be cases in which the claim being made (if it were being made as a legal claim) would be one of negligence, but it could include other kinds of claims as well. Sometimes serious personal injury is alleged to result from a breach of contract or promise or a trespass to property. In those cases, my argument would apply, and it would mean that if we were to find the alleged defendant causally responsible for these serious personal injuries, then he would necessarily be morally responsible too, even though no finding of traditional fault is required to make out a legal claim for either of these civil wrongs.

the natural end of life would be great, the incentives at work in evaluating whether there has been fault in this situation can be different: it is often comforting when we face the natural end of life of those we love to think that death was the inevitable result of a process in which human intervention was impossible. So we may choose not to pursue the matter further in these cases, even though if we did choose to do so the usual incentives that make proving fault relatively easy would apply to how *other* people evaluate what happened and we would therefore have no difficulty proving fault if we actually tried to do so. Nevertheless, a great deal of tort liability (indeed, almost all of tort liability other than what we would characterize as frivolous claims having only nuisance and not real monetary value) arises out of circumstances where people are seriously injured before the natural end of life. Regardless of whether the principles I discuss generalize to cases of less serious physical injury or to cases of economic injury alone, understanding the nature of causal inquiry in cases on which I have chosen to focus is critically important if we are to understand how the imposition of moral responsibility in many of the cases we are likely to confront actually works.

Note, however, that I will not be setting forth a comprehensive theory of causation in this paper, or attempting to analyze the countless difficulties that arise in the course of any causal inquiry. Those who are interested in an exhaustive survey of the concept of causation should accordingly look elsewhere.[5] What I am tackling in this paper is one very specific aspect of the relation between causal and moral responsibility, one that has surprisingly not been examined before despite the number and length of previous studies of causation: the extent to which assignments of causal responsibility constitute assignments of moral responsibility. Not whether

[5] See, e.g., Michael S. Moore, *Causation and Responsibility* (Oxford: Oxford University Press, 2009); H. L. A. Hart and Tony Honoré, *Causation in the Law* (Oxford: Oxford University Press, 2d ed. 1985); and J. L. Mackie, *The Cement of the Universe: A Study of Causation* (Oxford: Oxford University Press, 1980).

assignments of causal responsibility are *influenced* by moral considerations or involve the consideration of moral factors, but whether assignments of causal responsibility *are assignments of moral responsibility* full stop rather than something separate and apart from them. How certain difficult causal questions are to be resolved is accordingly irrelevant for purposes of this paper—all that is relevant is that we have some causal theory that conforms more or less roughly to our practice of assessing causal responsibility. It is that practice that forms the baseline for this paper, and it is this practice that I will claim is not merely a way of assigning causal responsibility but also a way of assigning moral responsibility, one that completely encompasses what we currently think of as being a separate and independent moral inquiry, whether inside or outside the domain of civil or criminal litigation.

Note also that I do not take any particular view in this paper as to what moral theory we should use to determine what constitutes a moral wrong (and therefore violates a moral right) or how serious we should consider that rights violation to be. Obviously, causal questions are often very important to these determinations whatever moral theory one adopts.[6] Indeed, even those who believe in collective moral responsibility must have a causal element somewhere in their theory.[7] Causal issues are often critical to assessing not only moral responsibility but moral seriousness too. My intent is not to comment on the nature of moral inquiry, but rather to re-position it within a larger and more important causal context. What I am interested in exploring

[6] I say "often" and not "always" here because some people think that causation is irrelevant to the determination of whether something constitutes a wrong to Immanuel Kant. Others think that this is a misinterpretation of Kant's theory. See generally Gardner, "Obligation and Outcomes in the Law of Torts," at pp. 141-142. Whether it is or it isn't, however, is irrelevant to anything I am discussing. The point is simply that even though causal issues may be relevant to determining whether something is a wrong, this paper is not an argument about how we should identify what wrongs there are. It is an argument about how we decide who is morally responsible for such wrongs no matter how they are identified.

[7] See Mark R. Reiff, "Terrorism, Retribution, and Collective Responsibility," *Social Theory and Practice* 34 (2008): 209-242, esp. 239-240.

is whether we should treat our causal and our moral inquiries as two separate and largely independent although perhaps somewhat overlapping exercises, as we do now, or whether the causal inquiry process already has embedded in it all the elements necessary to assign not only legal liability but moral responsibility as well, with or without a finding of traditional fault.

Of course, there is a long history to discussions of the relationship between causal and moral responsibility. Various theorists, among them those typically categorized as American Legal Realists as well as certain members of the Critical Legal Studies movement, have long claimed that there is no fact of the matter about causation—causation is simply the label we use to ascribe causal responsibility to a causal factor that was actually assigned responsibility on "policy" grounds.[8] And while it is never entirely clear what those who make this argument mean by "policy," it seems reasonable to interpret this as including although not necessarily being limited to a moral judgment (it could also be a prudential judgment). A similar "causal minimalist" position has also been taken at various times by those in the Law and Economics movement: Ronald Coase, for example, famously argued that we cannot determine what causes what until we have established a baseline determined on other grounds.[9] Whether the fire in a field of wheat was caused by sparks from a train or whether it was caused by the farmer planting his wheat too close to the tracks can only be resolved once we have decided who was entitled to do what.

My argument, however, is not a reprise of these views. I am *not* claiming that there is no metaphysical element to our causal determinations, only policy, however

[8] See, e.g., Wex Malone, "Ruminations on Cause-in-Fact," *Stanford Law Review* 9 (1956): 60-99; Mark Kelman, "The Necessary Myth of Objective Causation Judgments in Liberal Political Theory," *Chicago-Kent Law Review* 63 (1987):579-647.

[9] See, Ronald Coase, "The Problem of Social Cost," *Journal of Law and Economics* 3 (1960): 1-44.

policy may be defined. That contention I regard as having been convincingly rebutted long ago by Hart and Honoré, who showed that there were indeed some simple metaphysical principles regulating our common sense causal judgments.[10] Nor am I simply restating the Hart and Honoré thesis, for while they explained how metaphysical principles do indeed play an important part in our causal judgments, they also thought that additional factors beyond those necessary to come to causal judgments were necessary to come to moral judgments, and that causal responsibility was a necessary but not sufficient condition for moral responsibility, two propositions that I am going to deny.[11] What I am claiming is that all of the factors necessary and sufficient to make causal judgments are also necessary and sufficient to make moral judgments, and more precisely, once we have made the former, the actual content of the latter is effectively *entailed* by that causal judgment, at least in the subset of cases that I am going to specifically address. In other words, what we currently treat as a largely separate and independent moral inquiry to be conducted after our causal inquiry is complete is actually part of and subservient to what should be recognized as a comprehensive causal inquiry, one that does much more than merely set the stage for the real work to be done by whatever particular moral theory we happen to embrace. Rather than express a kind of causal skepticism, what I seek to do is to argue for a kind of causal all-inclusiveness, to take a look behind the curtain of causal inquiry and show that it actually amounts to an investigation of everything we need to know in order to assign moral responsibility too.

Now one of the implications of this position is that there can be moral responsibility for what in law are called strict liability offenses, given that we

[10] For an excellent description of the causal minimalist position and Hart and Honoré's rebuttal of it, see N. E. Simmonds, "The Dissolution of Law?," in *Where Next? Reflections on the Human Future*, ed. Duncan Poore (The Board of Trustees, Royal Botanic Gardens, Kew: 2000), pp. 209-227, 225-227.

[11] See Hart and Honoré, *Causation in the Law.*

typically (and correctly in my view) do not think that an attribution of traditional fault is necessary to establish causal responsibility. I will discuss this implication in some detail later in the paper, but right now I think it is important to make clear that while this is indeed an implication of my argument, it is not what my argument is all about. There could be moral responsibility for strict liability offenses even if my argument were not correct, so while my argument is sufficient for establishing this proposition it is not necessary—other arguments could establish this proposition just as well (indeed I think they do). And while nothing in the argument I am going to make here depends on the these other arguments being correct, the purpose of my argument here is not simply to add another arrow to this quiver. I consider the point that there can be moral responsibility for strict liability offenses relatively trivial, even though I suppose it remains somewhat controversial.[12] I am arguing for something much more significant than that. I am arguing that it is wrong to think of our ultimate moral judgment—the assignment of moral responsibility—as a higher-order judgment to be made only after we have made a series of causal judgments, and that the purpose of this ultimate moral judgment is to pick and choose among various causal candidates, only one of which will make the final cut. This of course tells us something important about strict liability cases, but it tells us something even more important about cases that seem to be based on findings of traditional fault, like ordinary negligence cases—it tells us we have been kidding ourselves by thinking that the findings of traditional fault that are made in these cases play the pivotal role we assume they do in determining whether the legal wrong of negligence is also a moral wrong. Of course,

[12] Especially, of course, because there are some big names taking the position that assigning strict liability is inconsistent with the idea of assigning moral responsibility: see, e.g., H. L. A. Hart, *The Concept of Law* (Oxford: Oxford University Press, 2d ed., 1961. 1994), pp. 166, 173, 178-179 (suggesting that there can be no ascription of moral responsibility without what I call traditional fault); Thomas Nagel, "Moral Luck," in *Mortal Questions* (Cambridge: Cambridge University Press, 1979), pp. 24-38, 31: "strict liability . . . may have its legal uses but seems irrational as a moral position."

the presence of traditional fault here *matters* to this determination, but my argument is that it is actually the assignment of causal responsibility that makes the defendants in these cases morally responsible for the injuries the plaintiffs have suffered, and that while the presence of traditional fault is *relevant* to the determination of whether causal responsibility exists in these cases, it is neither a necessary nor a sufficient condition for making such a determination. Instead of relying on the presence of traditional fault to drive our moral reasoning, what I intend to argue is that causal reasoning *is* moral reasoning, and the idea that there is separate process we must go through to arrive at an assignment of moral responsibility once we have made an assignment of causal responsibility is just an illusion. In other words, what I will have to say tells us far more about non-strict liability cases than it does about strict liability cases, although it tells us something about the latter too.

But even with regard to non-strict liability cases, my argument is intended to generalize rather than be limited to special problem cases. For example, my argument, if successful, would solve the problem that some people see in assigning moral responsibility in everyday negligence cases when we hold people to an objective standard of conduct (the degree of care we can expect to be exercised by the proverbial reasonable person) that, because of their inherent physical or mental limitations, they are not able to meet. It would also solve the problem that some people see in attaching greater moral blame to conduct that has severe consequences than to conduct that has no severe consequences or perhaps even no consequences at all when the conduct involved is the same, a problem that is often described as the problem of moral luck. But even these are side-effects of my argument, not its main thrust, which is that causal responsibility plays a different role in our judgments of

moral responsibility than we think it does, at least where there is serious physical injury. Not only in problematic cases, but in unproblematic cases too.

Finally, I want to emphasize that the point I am making in this paper is not simply that there is a moral element to every determination of causal responsibility. Again, I take this proposition to be relatively uncontroversial, for reasons which I will explain in detail in a moment. The conclusion that there is a moral element to every causal determination will of course be supported by my argument, but this is merely a preliminary point in the overall point I will be making. My overall argument is that casual judgments are themselves not merely moral judgments *in some sense*—these causal judgments do not merely have a moral *element* to them—they are actually assignments of moral responsibility. In other words, in the class of cases on which I will be focusing in this paper, everything we need to know in order to assign moral responsibility is there once an assignment of causal responsibility has been reached.

Something like this is behind the legal doctrine of *res ipsa loquitur* ("the thing speaks for itself").[13] The idea behind this doctrine is that in some situations, what I call traditional fault will be presumed, and so it is up to the defendant to rebut it rather than to the plaintiff to prove it. For example: even though a plaintiff who was sickened when he discovered a human toe in a package of tobacco could not prove how the toe got there and therefore that it got there because the packager was negligent, we can presume that human toes don't get into packages of tobacco unless the packager failed to exercise reasonable care.[14] But while my argument shares some of the intuitions behind the *res ipsa* doctrine (that someone must be at fault in certain situations even if we cannot prove intent or at least lack of reasonable care in

[13] See generally Restatement (Second) of Torts § 328D (1965).

[14] See Pillars v. R. J. Reynolds Tobacco Co., 78 So. 365 (Sup. Ct. Miss. 1918). See also Byrne v. Boadle, 2 H. & C. 722, 159 Eng. Rep. 299 (1863), where a barrel of flour rolled out of the window of the defendant's warehouse and fell on a passing pedestrian.

the ordinary way), my argument is not merely an attempt extend this to all cases of serious physical injury. I am not arguing merely for expansion of an evidentiary rule—I am arguing for something much more fundamental than this. I am arguing that in certain situations (situations in which serious physical injury has occurred), those who are causally responsible are morally responsible too, that this is true as a matter of both our legal and our moral practice although we currently fail to recognize this, and that while the presence of traditional fault is relevant to our assignment of causal responsibility, our conclusion that traditional fault is present is sometimes merely a *product* of an assignment of causal responsibility—in other words, we can find causal responsibility without traditional fault as an input, and when we do this traditional fault then appears as an output. So while traditional fault (in the form of at least unreasonable conduct) is present in every case in which someone is causally responsible, it is not true that this person is morally responsible *because* traditional fault is present, at least this is not true in every case. In some cases—that is, in cases of serious physical injury, an assignment of causal responsibility can be driven by other factors, and when this happens this assignment of causal responsibility effectively generates an assignment of moral responsibility all by itself. Indeed, given the nature of our assignments of causal responsibility, it is my view that our practice of assigning moral responsibility cannot be adequately explained in any other way.

Finally, I also want to make clear that my argument is not simply another entry in the ongoing debate over whether the idea of moral responsibility for outcomes and not just actions and omissions is intelligible.[15] This debate is one way the debate over the morally intelligibility of the imposition of strict legal liability is expressed,

[15] See, e.g., Tony Honoré, "Responsibility and Luck: The Moral Basis of Strict Liability," *Law Quarterly Review* 104 (1988): 530-553; John Gardner, "Obligations and Outcomes in the Law of Torts," in *Relating to Responsibility*, ed. Peter Cane and John Gardner (Oxford: Hart Publishing, 2001), pp. 111-144.

although the debate over outcome responsibility is actually wider than that for it also asks whether the imposition of legal liability for negligence under an objective standard can be morally intelligible. But as I hope I have already made clear, while my argument supports moral intelligibility here, this is a side-effect rather than the point of my argument. Indeed, if anything, my argument suggests that this way of framing of the question of whether there can moral responsibility in strict liability and objective standard negligence cases is incorrect, or at least misleading, because by doing so we are effectively assuming that moral responsibility is something to be decided after causal responsibility has been assessed. After all, to decide that one state of affairs is "an outcome" of any particular individual or entity's conduct means one has already determined that a particular actor "caused" that outcome in the relevant sense. This way of phrasing the issue therefore begs or at least conceals the question I am addressing in this paper; that is, whether an assignment of causal responsibility already has within it an assignment of moral responsibility too. It is to the task of exploring this question that I now turn.

II. WHAT IS CAUSAL RESPONSIBILITY?

The first issue I want to address is how we are to understand causal responsibility. For causal responsibility is usually split into two parts: *actual* causal responsibility and *proximate* or *legal* causal responsibility.[16] Determining actual causation is supposed to pose a purely empirical question about how one historical event relates to another. Determining legal causation, in turn, is not really meant to be a causal inquiry at all, but rather an “all-things-considered” enterprise in which various policy reasons for holding the defendant liable are balanced against various

[16] I will use the term "legal cause" to refer to this kind causation from here forward, but this term should be understood as having the same meaning as "proximate cause" for those more familiar with that term.

policy reasons for not doing so.[17] As a policy matter, the question of legal causation is accordingly already understood to be a form of moral inquiry, one that raises questions of distributive if not commutative justice; one that is informed by larger empirical questions about the social consequences of the various alternatives open to us but one that is fundamentally a question about justice nonetheless.

To illustrate this, let me describe a situation that is handled differently in England than it is in California. Suppose a drunk driver runs over a pedestrian, seriously injuring but not killing him. Unfortunately, the driver is uninsured and has no assets from which any judgment against him might be paid. But the owner of the bar that served the driver has substantial assets, and in any event is insured. Can you successfully bring a claim for negligence against the owner, arguing that he (through his employees) knowingly served someone who was inebriated and then allowed him to drive away (perhaps the bartender was asked for directions by the man who was too drunk to remember how to drive home). Interestingly, the answer is that you can make such a claim in California, and you can indeed recover if you can prove that your claim is true, but you could never get that far with such a claim in England.[18] But why should this be so? The causal connection of the bar owner to the unfortunate incident is the same in both countries, both jurisdictions allow claims for negligence, so why the different result?

[17] See, e.g., Lord Denning's remarks in *Lamb v Camden London Borough Council* [1981] QB 625, at 636.

[18] See, e.g., Vesely v. Sager, 5 Cal.3d 153, 486 P.2d 151 (Cal. 1971) (California); Barrett v. Ministry of Defence [1995] 1 WLR 1217 (England). Note that in Barrett, the question addressed was phrased as one of duty of care, not causation. But all questions of causation can be phrased as questions of duty of care, and vice versa—this is ultimately just a matter of semantics, although some particularly stubborn academics pretend that it is not. In any event, while Barrett does not absolutely resolve the question of whether those who provide alcohol to others can be liable to those these others then injure in some way, I think any fair-minded English lawyer would consider such a claim a long shot, while any fair-minded California lawyer would not.

At this point, the explanation typically offered by my students is that people in California are crazy, whereas people in England are far more sensible—people in England can see that the owner of the bar is far too removed from the actual act of wrongdoing in this case to be seen as a proximate cause of it. As appealing as this answer may be on a superficial level, however, I think that something deeper and more revealing is actually going on. The thinking behind the results in these two jurisdictions is actually the same; the result differs merely because the contingent factual context of the accident is different in each jurisdiction. If the drunk driver is uninsured, the injured plaintiff will be left without a remedy in California, and if that plaintiff has no medical insurance of his own, he could be left with no way to pay for the medical attention he now needs. But this would not be the case in England. In England, if the driver is uninsured, the injured plaintiff can always recover from the Motor Insurers' Bureau ("MIB"), and in any event he can always obtain the same medical care he could obtain even if the driver was insured through the National Health Service, or NHS. Because the injured party is not left without a remedy in England, and is not faced with having to pay for his medical care on his own, there is no need to impose liability farther back along the chain of causation in such cases.

But California has no equivalent of the MIB, and no equivalent of the NHS either. So in California, distributive justice provides a much stronger reason to find someone farther back along the chain of causation that ultimately led to the plaintiff's injury responsible for it. Doing so allows us to redistribute the burden of the plaintiff's injury to a wider pool of individuals through the bar owner's insurer. It is therefore driven by a concern for distributive justice, and this outweighs any commutative justice concern that this more remote defendant might be unfairly surprised by the imposition of liability on him the first time this actually occurs.

What this comparative example shows, therefore, is how a determination of legal or proximate cause can be influenced by moral considerations, rather than being ruled exclusively by physical and metaphysical ones. Of course, one example illustrates but does not prove the case. There are enough other examples set forth in the existing literature, however, that I hope it is unnecessary to trot out more here.[19]

Of course, while the principles of legal causation can sometimes be used to reach farther back along the chain of causation to impose legal liability (I would say moral responsibility but I have not shown this yet so legal liability will have to do for now), legal causation is usually seen as a limiting principle—the idea of actual causation throws a very wide net, bringing in all sorts of people who are in some sense causal contributors to the events at issue and therefore are potential defendants in a lawsuit. The concept of legal causation is then used to whittle down this very large number of potential defendants to just a few, for it would simply be unmanageable if not unfair to hold everyone who is a causal contributor legally liable given their potential remoteness from the event at issue.

But sometimes it seems that the net thrown by actual causation is not wide enough. In some cases, the traditional "but for" test of legal liability—the idea that someone is an actual cause of an injury if the injury would not have happened "but for" his conduct—excludes certain people as potential defendants that our moral intuitions tell us should actually be included. So while no court has yet officially adopted it as a test in any country, for purposes of assigning moral responsibility rather that legal liability I will treat the actual causation question as being determined by what Richard Wright calls the "NESS" test. The NESS test provides that one

[19] For example, for further discussion of the various ways that moral issues—or, as they are sometimes euphemistically referred to, "policy" questions—can enter into the proximate cause inquiry, see Moore, *Causation and Responsibility*, pp. 89-104.

event is the actual cause of another if and only if it was a "necessary element of a set of conditions that was sufficient for the occurrence of the result." This test is based on an account of causation first elaborated by David Hume and subsequently modified by J.S. Mill. As Wright describes it

> Hume revolutionized philosophical thinking on causation when he insisted that, contrary to then-popular belief, causal judgments are not based on the direct perception of causal properties or forces inherent in objects or events, [but rather] on the belief that a certain succession of events instantiates one or more causal laws. A fully described causal law would list all of the conditions that together are sufficient for the occurrence of a certain consequence. Yet, to avoid including causally irrelevant conditions in this sufficient set, the antecedent conditions are restricted to those that are necessary for the sufficiency of the set. Thus the necessity requirement is subordinate to the sufficiency requirement. When the foregoing is combined with Mill's observation that there may be a plurality of distinct sets of conditions that are each sufficient to produce the consequence, so that there is no uniquely sufficient set, the NESS test is complete.[20]

Unfortunately, however, there are cases in which even the NESS test does not ensure that we have included everyone in the net of potential defendants that we might want to include. These cases are colloquially referred to as the problem of the indeterminate defendant, the problem of the indeterminate plaintiff, and the problem of the indeterminate cause. In the first case, you know that one of several defendants caused the particular injury, but you do not know which one. In the second case, you know a particular defendant caused injury to say, 45% of the people who have come down with a certain disease, but you do not know if the particular plaintiff is one of these people or someone whose disease had some other cause. And in the last case, you know there is a 45% probability that the defendant's actions caused your injury, but your injuries were more likely than not the result of some another cause. In each

[20] Richard Wright, "Causation, Responsibility, Risk, Probability, Naked Statistics, and Proof: Pruning the Bramble Bush by Clarifying the Concepts," *Iowa Law Review* 73 (1988): 1001-1077, at 1019-20 (footnote omitted). See also Richard Wright, "Causation in Tort Law," *California Law Review* 73 (1985): 1735-1828; and "Once More Into the Bramble Bush: Duty, Causal Contribution, and the Extent of Legal Responsibility," *Vanderbilt Law Review* 54 (1995): 1071-1132.

of these cases, at least some people's intuitions suggest that we have reason to hold each of the potential defendants legally liable for some portion of the plaintiff's injury (and partially morally responsible too), even though we cannot be sure that their actions were an actual cause.[21] So in many cases, there is good reason to think that even when we are talking about actual causation, we need to make a moral judgment in the course of deciding whether we can impose causal responsibility and therefore legal liability as well.[22]

But as I said when I began, this is only a preliminary step in my argument. Proving that there can be a moral element to determinations of both actual and legal causation does not prove that there is *always* a moral element in every causal determination, much less that this moral element, whatever it may be, is sufficient for an ultimate assignment of moral responsibility. To take the next step in that argument, I need to explore the relationship between legal liability, which is what I have been talking about so far, and moral responsibility in much greater detail. I have to show that our judgments of legal liability are not simply judgments that may or may not map on to judgments about morality—I have to show that judgments about legal liability, at least in cases of serious personal injury, are moral judgments too. It is to that task, accordingly, that I now turn.

III. LEGAL LIABILITY AND MORAL RESPONSIBILITY

[21] See, e.g., *Sindell v. Abbott Labs*, 607 P.2d 924 (Cal. 1980), which involved an indeterminate defendant (the plaintiff was allowed to recover a portion of her damages from each maker of the drug DES according to that defendant's market share of sales of the drug at the relevant time, even though plaintiff could not prove which defendant had been the source of the drug that her mother had taken and which had caused her injury); Richard Delgado, "Beyond *Sindell:* Relaxation of Cause-In-Fact Rules for Indeterminate Plaintiffs," *California Law Review* 70 (1982): 881-908 (arguing for an extension of the *Sindell* approach to cases of where the identity of the injured plaintiff is indeterminate); David A. Fischer, "Tort Recovery for Loss of a Chance," *Wake Forrest Law Review* 36 (2001): 605-655 (discussing cases that have allowed recovery where the probability that defendant's act caused plaintiff's injury was less than 50% and advocating wider application of the doctrine).

[22] There are so many theorists that now agree with this, see, e.g., Jane Stapleton, "Cause in Fact and the Scope of Liability for Consequences," *Law Quarterly Review* 119 (2003): 388-425, I think it is fair to regard this proposition as non-controversial.

The reason why this is the next step in the argument is that everything I have said about moral responsibility so far could simply be re-characterized as a discussion of legal liability. People are held legally liable all the time, the argument goes, without necessarily being morally responsible. There may indeed be all-thing-considered reasons to hold people legally liable, and therefore liable to pay compensation, and some of these may be moral reasons, but some of them may be simply prudential, which means that those held legally liable are not necessarily also morally responsible. The latter move has yet to be justified, and if it cannot be justified then it seems I cannot move forward with my argument.

The first thing to remember here once again is that I am not arguing that an assignment of legal liability constitutes an assignment of moral responsibility in all cases—on the contrary, my point is limited to cases where legal liability is assigned for serious personal injury. And as I have already noted, the reasons that go into the determination that we should shift the financial burden of a serious injury from the victim to someone else often are the results of a moral inquiry, for this turns on questions of distributive justice. We could, of course, decide to shift this financial burden from the victim to the taxpayer, as they do in New Zealand, for example, rather than to a particular entity or individual, but even this decision is not produced solely by considerations of economic efficiency—it is an expression of distributive justice concerns too, that is, our desire to ensure that the financial burdens of serious physical injury are distributed throughout society rather than allowed to fall on the unfortunate victim or even the negligent wrongdoer and his insurer alone. Even cases of vicarious and strict product liability contain elements of moral inquiry. We are vicariously liable for the torts of those under our supervision and control because there is something about our relationship to these individuals that differentiates us

from the general population, for the nature of this relationship puts us in a position to take action to see that those under our control behave appropriately, even when we are not there to supervise their every move. Similarly, with regard to strict products liability, manufacturers are not *absolutely* liable for any injuries their products happen to cause; they are liable only if that product was *unreasonably* dangerous, so even this form of supposedly fault-free liability has a moral notion embedded in it.[23]

For those that find these embedded moral notions insufficient to change the causal element within legal liability for serious physical injury into a determination of moral responsibility, however, consider this: tort law, no matter how popular the notion, is not simply about compensation. It is about punishment too. Indeed, if it were not also about punishment, it is hard to see how tort law could ever act as a deterrent to wrongful or at least dangerous conduct, as it is designed to do. The obligation to exercise reasonable care is designed in part to encourage people to take measures to avoid injuring their fellows; the existence of vicarious liability is meant to encourage supervision and training of those under one's control; the imposition of strict products liability is not only designed to ensure compensation but also to ensure that products put into the stream of commerce are safer than they would otherwise be. None of this could be accomplished if the imposition of legal liability did not threaten punishment as well as provide a means for obtaining compensation.

True, a great deal of legal liability is covered by insurance. This means that the burden of paying compensation is often shifted from the party causing the injury to some other party, and therefore compensation rarely comes from the actual wrongdoer's pocket. But that does not mean that wrongdoer is able to fully shift the

[23] See Restatement (Second) of Torts § 402A(1) (1965). Article 6(1) of EEC Directive 85/374 incorporates a similar provision into EU law by requiring that the product is defective only if it does not provide the safety which a person is entitled to expect.

consequences of legal liability to another. Some consequences remain, and these consequences qualify as punishment, which is why legal liability acts as a deterrent and not merely as a means of shifting the financial burdens associated with serious injury. Doctors vigorously fight malpractice claims (even the settlement of malpractice claims) because they are concerned that anything other than an outright dismissal of the claim is going to damage their reputation, notwithstanding the fact that any payment out is going to be covered by insurance. Manufacturers apologize for defects in their products (for a recent example, think of the lengths Toyota has gone to apologize for the defects recently discovered in its cars and assure the public that such defects will not arise in future models) because they fear damage to their reputation and the financial loss that this can cause, even though any direct financial payments out may be covered by insurance.[24] And so on.[25] Indeed, if all the consequences of legal liability could be shifted to other parties then it is hard to see how legal liability could act as a deterrent at all. So civil wrongs have to trigger some punishment too—that is, "an undesirable change in a person's well-being that is seen to result from or is deliberately imposed as a response to some act or omission committed by that person or by others for whom he is deemed responsible" that cannot be shifted to others.[26] Finally, remember that punishment is traditionally seen

[24] See Micheline Maynard and Martin Fackler, "Repairing Some Dents in an Image," *The New York Times* (August 5, 2006): Hiroko Tabuchi, "1.5 Million Toyotas Recalled for Brake and Fuel Pump Problems," *The New York Times* (October 21, 2010); Hiroko Tabuchi, "Toyota Recalls 1.7 Million Vehicles," *The New York Times* (January 26, 2011).

[25] For further discussion of the natural limits on our ability to shift the burden of the consequences of liability to others, see Mark R. Reiff, *Punishment, Compensation, and Law: A Theory of Enforceability* (Cambridge: Cambridge University Press, 2005), pp. 215-221.

[26] I realize that some people deny that tort law is about punishment at all, especially those who seek to describe tort law as an instantiation of pure corrective justice. But I do not think this view is defensible for the reasons I have set forth in the text. For those who remain unconvinced, however, more extensive argument on this point is set forth in my *Punishment, Compensation, and Law: A Theory of Enforceability* at pages 75-98. The definition of punishment that I rely on in the text is found and defended on page 77 of that volume.

as unjustified in the absence of a corresponding degree of moral responsibility.[27] Accordingly, we must either accept that the entire tort system is hideously unjust, or recognize that legal liability—even strict legal liability—is necessarily a form of moral responsibility.

What this means is that assigning legal liability is not merely a moral decision in the distributive sense. True, part of the decision to assign legal liability may based on distributive concerns—we are trying to ensure that the plaintiff has access to compensation so that the cost of his injury can be distributed over a larger group of people, either through one party's insurer or the other's or through the government to the taxpayer. But there is also something more personal going on in these cases, something beyond a desire to ensure access to compensation, for the cost of the injury is not the only thing that is being dealt with. As I have argued at length elsewhere, most injuries, and I think it is safe to say that all serious injuries, involve some incompensable injury.[28] So deciding where legal liability lies does indeed do some moral work. It tells us how to distributive moral praise and blame, and who is entitled to and who is subject to punishment and retribution. So when we are assigning causal responsibility in these cases, we are doing more than making a possibly impersonal but in any event distributive moral judgment. We are also making a personal commutative moral judgment, because we are not merely deciding how the plaintiff is to be compensated for his injuries, we are also deciding whether this particular defendant—this particular causal factor—deserves to be punished for what he or she or it has done.

[27] See H.L.A. Hart, *The Concept of Law* (Oxford: Clarendon Press 1961), pp. 168–76, esp. p. 173.
[28] See Reiff, *Punishment, Compensation, and Law*, pp. 215-221.

I recognize, of course, that this is contrary to the conventional view. The conventional view is that wrongs should be divided into two categories: the civil and the criminal. Civil wrongs are to be compensated, even if they are accompanied by proof of fault; only criminal wrongs are to be punished. This, in turn, makes deciding whether punishment is due dependent on whether the wrong at issue is properly characterized as civil or criminal, and this in turn is usually seen as dependent on whether the wrong in question is sufficiently serious in terms of its broader societal effects to be the public's business, and not merely a private matter. In response to this view, some theorists argue that the test should not be whether the wrong has significant societal effects, for many of the wrongs we currently accept as civil clearly have such effects and many wrongs we currently accept as criminal do not, or at least do not always have such effects. On the contrary, the real dividing line between the civil and the criminal—between whether compensation or punishment is due—should be determined by the intrinsic nature of the wrong itself. Of course, in deciding what intrinsic factors make a wrong more criminal than civil and what factors have the opposite effect, the same factors may come into play, so the real difference between these two views may merely be semantic. In either case, however, we are looking for some criterion that establishes the difference between the civil and the criminal.[29]

But there is another possibility here. Instead of focusing on whether a wrong should be classified as civil or criminal in order to determine whether it is punishment or compensation that is due, we might recognize that punishment as well as compensation can be due for any wrong, regardless of whether it is civil or criminal. This is the argument I present in *Punishment, Compensation, and Law: A Theory of*

[29] See generally Antony Duff, "Towards a Theory of the Criminal Law?" *Aristotelian Society Supplementary Volume* 84 (2010): 1-28; Douglas Husak, *Overcriminalization: The Limits of the Criminal Law* (Oxford: Oxford University Press, 2008).

Enforceability, and while I cannot adequately summarize that book-length argument here in just a paragraph or two, I can sketch out its basic structure. What I argue there is that what determines whether punishment is due is not that nature of the wrong, but whether and if so to what extent the wrong causes incompensable suffering. To the extent a wrong causes merely compensable suffering, then that suffering should be compensated, for punishment is the remedy of last resort. To the extent it causes suffering that can't or won't be compensated, however, then punishment is required to restore the moral balance between the wrongdoer and the wronged. There is almost always *some* kind of punishment, moreover, that can be meted out no matter how we characterize a wrong. For example, the issuance of moral criticism, the experience of moral regret, and certain forms of social sanctions are forms of punishment that will almost always be available. But in some cases the nature and extent of the incompensable suffering caused by the wrong will be such that only the kinds of sanctions available under the criminal law (that is, criminal fines and imprisonment and the more severe kinds of social sanctions that arise from fines and imprisonment) will be sufficient to put the post-violation moral balance right. It is in these cases, then, where the wrong should be classified as not merely civil but also criminal.[30]

Of course, we still need to decide what kind of acts are wrongs and what are not, and this may indeed depend on the intrinsic nature of the act in question or its societal effects or both, for not all acts that cause incompensable suffering are wrongs, but this has nothing to do with whether compensation or punishment is due. While the latter question is *determined by* the need to enforce whatever rights we have, it is not a factor in *determining* what rights we have. So the fact that we do punish people

[30] For the full version of this argument, see Reiff, *Punishment, Compensation, and Law: A Theory of Enforceability.*

even when a finding of moral fault is not an official element of the wrong does not mean that we impose punishment without regard to fault. On the contrary, it means that we may have found moral responsibility without regard to *traditional* notions of fault, but that any case in which we impose legal liability at least purports to be a case where we can assign moral responsibility. And no matter whether we have found traditional fault or not, the moral responsibility that we have found flows simply from our finding causal responsibility. How we go about finding causal responsibility is accordingly what I will attempt to explain next.

IV. CAUSAL CONTRIBUTION AND CAUSAL RESPONSIBILITY

As Bernard Williams noted long ago, a person who bears causal responsibility for a serious injury bears a special relationship to that injury, a relationship that is very different than the one enjoyed by those who are in no way causally responsible for the injury.[31] The example Williams gives is that of a lorry driver who has run down a small child. In Williams' example, the lorry driver is not at fault Williams says, but nevertheless feels differently than mere spectators to the incident and, Williams argues, *should* feel different. What Williams does not specify is whether the accident was solely the child's fault or no one's fault. To the extent that he has the latter possibility in mind, of course, I disagree that this could be the case—that is, I disagree that no one could be at fault. To the extent that Williams has the former possibility in mind, however, we have to ask whether he could be correct. Can the lorry driver be causally responsible but not morally responsible? Is this why Williams and the rest of us think he should feel different than a mere spectator?

[31] See Bernard Williams, "Moral Luck," in *Moral Luck* (Cambridge: Cambridge University Press, 1981), pp. 20-39, at pp. 27-28. Similar sentiments have been expressed by others as well. See, e.g., Moore, *Causation and Responsibility*, p. 30.

Of course, in answering these questions, we have to specify what we mean when we say the lorry driver "feels different." The only way to do that, in my view, is to say that the lorry driver feels some sort of normative regret. He feels he should have done something other than he did, and that "should," in turn, makes most sense as a moral should. I suppose it could be a sort of prudential should—the kind of should we mean when we think "I should not have eaten so much at dinner." But it seems odd to think that this is all we mean when we say the lorry driver should feel different in Williams' example. The "should" we have in mind seems to have much more gravitas than this. And for the lorry driver to rationally feel moral regret, he must not be a mere causal factor—that is just one of the hundreds or thousands or perhaps even millions of previous acts and omissions by large numbers of people that were in some sense necessary for this truck and this boy to end up at the same place at the same time. No, to feel different in the way Williams is suggesting, it must be the case that he feels morally responsible, and if he feels morally responsible he must feel causally responsible too. Not, perhaps, in the same way and to the extent as someone who failed to exercise reasonable care, but causally and morally responsible nonetheless. His responsibility may be dwarfed in comparison to that of someone else (the child itself, for example), and therefore we may not hold him legally liable, but the only way the lorry driver would bear no moral responsibility at all is if he bore no causal responsibility either.

Now this could be the case—it could be that running over the child involved no voluntary act on his part (perhaps the truck was being operated by someone else by remote control and the driver had no way to communicate with his person or take the wheel of the truck himself), although this seems unlikely. In that case, the lorry driver would bear neither causal nor moral responsibility. But even if the child darted out in

front of him in an area where children do not usually play, it seems more likely that we would find him 5 or 10 percent responsible, but part responsible nonetheless, for it would not be true that he was a mere involuntary instrument in a causal chain in which he otherwise had no involvement or control. And if the incident did indeed involve a voluntary act by him that we are willing to treat as causally responsible for what happened, at least in part, then we would indeed consider him morally responsible, and he would feel regret and indeed should because he does indeed have something to regret. Which means that to the extent that Williams is trying to suggest the contrary, his example is misleading, for it does not accurately represent our common causal and moral judgments. We simply could not treat the lorry driver any differently than a mere spectator (who could also no doubt have interfered in some way and tried to save the child) unless we think that the lorry driver is *both* causally responsible *and* morally responsible, at least in part.

Some support for this phenomenon can be found in the empirical evidence surrounding what has come to be known as the "Knobe effect." In brief, Knobe found that people generally hold others both causally and morally responsible for *harmful* side-effects that are foreseen (or perhaps merely foreseeable) even if not expressly intended, but not causally or morally responsible for *beneficial* side-effects under the same circumstances.[32] This asymmetry seems to be present in other circumstances too.[33] Exactly what lessons we should draw from this, however, are somewhat unclear—Knobe claims that among other things these results show that our moral views shape our causal judgments, but it is equally possible that they show our causal views shape our moral judgments or, as I argue here, that we are really only making

[32] See, e.g., Joshua Knobe, "The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology," *Philosophical Studies* 130 (2006): 203-231.

[33] See Fiery Cushman, Joshua Knobe, and Walter Sinnott-Armstrong, "Moral Appraisals Affect Doing/Allowing Judgments," *Cognition* 108 (2008):281-289, 285-288.

one judgment and it has both causal and moral elements to it. Indeed, my explanation for why we are so ready to find both causal and moral responsibility for serious personal injury explains the asymmetry Knobe found: we find causal and moral responsibility because we don't want to believe we live in a dangerous world, but we do not find responsibility for benefits under the same circumstances because we like the idea that good things can "just happen." In any event, Knobe's experiments at the very least support my claim that our causal and moral judgments rise and fall together.[34] And while this does not prove that the process of causal inquiry does involve the assessment of everything necessary to assign moral responsibility too, this does seem to be the best explanation we have available for this.

Of course, one could say that all the asymmetry discovered by Knobe really shows is that there is an inconsistency and therefore some irrationality in our common reasoning that we need to be aware of and overcome—the presence of this inconsistency need not be taken to suggest that this is the way either causal or moral inquiry should be conducted. Indeed, some scholars seem inclined to dismiss this aspect of our practice for this very reason.[35] But I disagree that there is anything irrational about this asymmetry, or anything inconsistent here that should be corrected. I am not even sure where such a "should" would come from. Moral judgments are reactive attitudes, and it would be odd indeed if we were to define something as morally good even if we did not think it was and as not morally wrong even if we did think it was. Morality is not independent of common features of human nature and our reasoning abilities, it is an expression of them. Although morality need not always coincide with our moral intuitions, it would be hard to know

[34] For similar experimental findings, see Mark D. Alicke, "Culpable Control and the Psychology of Blame," Psychological Bulletin 126 (2000): 556-574.

[35] See., e.g., Rachlinski, "Misunderstanding Ability, Misallocating Responsibility."

what morality was if it did not accord with our moral intuitions in large part. There are certainly consequential arguments that could be made in favor of this asymmetry (arguably, we are all better off if we are relentless in our pursuit of how to prevent injury, no matter what precautions have been taken, but not so ready to reward what amounts to the unintended conferral of benefits). In any event, assigning causal and moral responsibility is a human practice, and it seems both unrealistic and unhelpful for us to insist that any human practice ignore what most of us see are real differences between causing harm and simply being around when a foreseeable benefit happens to occur.

Nor is this a misguided attempt to derive an "ought" from an "is." The claim that it is not possible to do this is usually attributed to Hume, although Hume actually said something much more equivocal—that we simply have to be careful when it appears we are moving from descriptive statements to evaluative conclusions.[36] Many theorists actually disagree with Hume's view, or rather the categorical version that is typically attributed to Hume.[37] And it does seem incorrect to think that if morality is built upon our common reactive attitudes, or at least on the common reactive attitudes of rational people, one cannot use those reactive attitudes as least as some sort of evidence (even if not as conclusive evidence) that a particular form of

[36] See David Hume, *A Treatise on Human Nature*, ed. L. A. Selby-Bigge (Oxford, 1896 ed.), bk. 3, pt. 1, sec. 1, p 469: "In every system of morality, which I have hitherto met with, I have always remark'd, that the author proceeds for some time in the ordinary ways of reasoning, and establishes the being of a God, or makes observations concerning human affairs; when of a sudden I am surpriz'd to find, that instead of the usual copulations of propositions, *is*, and *is not*, I meet with no proposition that is not connected with an *ought*, or an *ought not*. This change is imperceptible; but is however, of the last consequence. For as this *ought*, or *ought not*, expresses some new relation or affirmation, 'tis necessary that it shou'd be observ'd and explain'd; and at the same time that a reason should be given; for what seems altogether inconceivable, how this new relation can be a deduction from others, which are entirely different from it.

[37] For some examples, see W. D. Hudson, ed., *The Is Ought Question* (London: Macmillan, 1969).

conduct is a moral wrong.[38] But even if the categorical version of Hume's view is correct, this is not what I am doing here. I am not deriving an "ought" statement from a series of descriptive "is" statements about people's reactive attitudes. I am arguing that people ought not to "cause" serious injury, and that "cause" is to be understood not merely as a descriptive term, although it surely has a descriptive element, but a term that has a moral, evaluative element built into it too even if it does not include what we think of as a finding of traditional fault.

To make this moralized version of causal responsibility work, however, we must distinguish between mere causal contribution and causal responsibility. For almost any event there are countless causal factors—that is, actions or omissions that lie somewhere on one branch or another of a causal stream that together with other causal streams create the chain of events that ultimately produces the particular injury in question. There are even more counterfactual causal factors—that is, acts that if undertaken might have led to a different result. Indeed, there are probably an infinite number of such factors.[39] It is pointless to deny that all of these actions and omissions make a causal contribution to every injury. Nevertheless, very few of them will be candidates for being assigned causal responsibility. While many of these factors may be causal in a strictly physical or metaphysical sense, they will not be causal in the ultimate assignment of causal responsibility sense. And many of the causal factors that *are* candidates for being assigned causal responsibility will not ultimately be

[38] Note that this does not amount to a claim that moral judgments cannot exist independently of our recognition of them—mine is not an attack on moral realism. I am simply arguing that causal judgments *are* moral judgments, even though they may be made without a finding of traditional fault. The metaethical status of these causal judgments is accordingly open to debate just as the metaethical status of traditional moral judgments would be. I am not saying that moral judgments must be recognized as such by (at least some) people to exist. I am merely saying that recognizing a judgment as a moral judgment is some evidence that it is.

[39] As David Lewis says: "Any particular event . . . stands at the end of a long and complicated causal history . . . the only question is whether [these causal chains] are infinite or merely enormous." David Lewis, "Causal Explanation," in *Philosophical Papers Volume 2* (Oxford: Oxford: University Press, 1987). pp. 214-240, at p. 214.

assigned such responsibility, for there will be intervening causes later on in the chain of actual or counterfactual causal factors that we decide are *supervening* causes, thereby trumping the causal responsibility of earlier causal factors. Or, alternatively, we may decide that causal factors that appear earlier in the chain makes such an overwhelming contribution to the ultimate event under consideration that we treat them as *overwhelming* causes that relieve even later causes factors of any casual responsibility.

Let us take an example to help us understand this idea of causal responsibility better. Suppose a baseball slugger who strikes out a lot asks his batting coach for some advice and his coach says "take fewer at-bats." We probably would find such advice comical, not serious. But of course in a technical sense the coach is absolutely right—surely taking at bats is a causal factor that influences the number of strike outs the slugger has. If he takes fewer at-bats, the number of his strike outs will almost certainly go down, other things being equal.[40] But this is not a causal factor that is of any interest to the slugger or that should be of interest to anyone trying to understand the problem the slugger is currently experiencing. Causal questions are always purposeful—a causal explanation is only meaningful if it is responsive to the purpose of the causal inquiry at hand.[41] In the case of the baseball slugger, what he is seeking to identify is whether there is something he can change about his current practice that will result in fewer strikeouts without negatively affecting what he has otherwise been

[40] To those not so familiar with the American game of baseball, an "at bat" is the opportunity a hitter has to hit the ball when it is his turn in the batting order to come to the plate to face the pitcher. A "strike out" occurs when the batter has three strikes. A strike, in turn, occurs when the batter swings and misses a pitched ball, or fails to swing at a pitch deemed in the strike zone, or fouls off a ball (hits the ball out of play) when he has less than two strikes on him already. Obviously, the fewer games the player plays, the fewer "at bats" he will have, and the fewer at bats he has, the fewer opportunities he will have to strike out.

[41] For a discussion of this feature of causal inquiry, see Hart and Honoré, *Causation in the Law*, pp. 32-44 (referring to what I have called causal factors as "conditions" and causal factors to which causal responsibility is assigned as "causes"). See also Lewis, "Causal Explanations," at pp. 226-227, listing the ways in which causal explanations can fail, especially example 5.

accomplishing at the plate. He wants to know if he is dropping his left shoulder too much, or gripping the bat wrong, or something of that nature. When we are dealing with the cause of serious injuries, in turn, what we are seeking to identify is whether there is something in the chain of causation that we can do something about so as to make such injuries less likely in the future, or at least more predictable and therefore more avoidable. We are looking for something that is under or at least could be put under human influence or control, a human act or omission that has in some significant way and to some significant extent contributed to the injury in question. There is accordingly a difference between identifying causal contributions and assigning causal responsibility—the former is simply a neutral act of historical accounting; the latter requires an evaluative stance. The latter is a subset of the former, but only the latter tells us what we want to know.

I therefore disagree with J. L. Mackie, who said, "[causation] is still a fairly unitary concept: we do not have one concept for physical causation and another for human actions and interactions."[42] That statement, while true in one sense is false in another: in theory, we *could* treat the contribution of human agents like any other causal factor, but in practice this would be pointless in most cases so we typically do not. In practice, we apply different theories for determining causation depending on what kind of causal inquiry we are undertaking. When we are asking what contribution, if any, human agents have made to a particular event, we are asking a very different question than we are asking when we ask what makes the earth revolve around the sun. It is the former type of question we are asking not the latter when we are inquiring into the cause of human misery, at least in all but the most unusual of cases. The former type of inquiry is a search for human causal responsibility, which

[42] Mackie, *The Cement of the Universe*, p. xi.

is an evaluative project by its very nature; the latter is a search for a purely non-evaluative, natural, scientific, descriptive explanation of the workings of the physical world.[43]

So how do we go about assigning causal responsibility? I have already gone over some of the factors that go into this inquiry—the kind of factors traditionally associated with the notion of "proximate" or "legal" cause. We balance our desire to ensure that plaintiffs are compensated for their injuries against our desire not to subject potential defendants to unfair surprise by imposing a moral and financial burden upon them they could not reasonably anticipate and therefore did not have an opportunity to insure against and take action to prevent. But there are other important factors here as well. The traditional notion of fault as a lack of reasonable care, of course, is one of them, but remember, when we are assigning causal responsibility, at least in the context that we are currently considering the matter, we are looking for more than simply unreasonable conduct—we are looking for guidance as to how we should react to any human action or inaction that ultimately makes a causal contribution to a serious human injury. All assignments of causal responsibility to such human action are accordingly morally significant, for they tell us how we might re-arrange the world to prevent a similar case of human injury from happening again.

This means we must recognize that the assignment of causal responsibility is in part a comparative inquiry, depending not only on the relation of a particular causal factor to an injury but also on the relation between that causal factor and the injury and other causal factors and the injury.[44] In other words, it is impossible to assign

[43] For an insightful and thorough discussion of this distinction, see Marion Smiley, *Moral Responsibility and the Boundaries of Community* (University of Chicago Press, 1992), ch. 8, esp. pp. 178-195.

[44] For a similar idea, see Jonathan Schaffer, "Contrastive Causation and the Law," *Legal Theory* 16 (2010): 259-297, and the various works cited therein at 261 n. 6.

causal responsibility to a causal factor without looking at the entire causal chain, or at least the last part of that causal chain, that bit of the causal chain that arises like an iceberg out of an otherwise overwhelmingly infinite causal sea. Where there is only one human causal factor, no matter how small, that factor is potentially significant (and by human causal factor here I am referring to both human action and inaction). If it is also relatively close along the chain of causation that leads to the injury, we are very likely to assign causal responsibility to it. Indeed, even a very small human causal factor may bear causal and therefore moral responsibility if there is no other human causal factor to bear it. In contrast, where there are several human causal factors (as there almost always are), we must look to see if their causal contributions to the injury vary greatly or are roughly equal. If some causal factors contributed significantly to the injury, and some contributed minimally, we are likely to assign causal responsibility only to the significant factors and ignore the others, for it is rational to start with the more important causal factors first. If the contributions of all these human causal factors are similar, however, we are likely to assign causal responsibility to all of them no matter how minor each individual factor's causal contribution may actually be. And a causal factor that is closer to the end of the chain of causation will often be assigned more causal and therefore more moral responsibility than those causal factors further down the line because we are likely to be more certain about the contribution of these factors to the injury at issue, although metaphysical closeness and remoteness is only one consideration to be taken into account in this comparative inquiry, and sometimes more remote factors may be assigned greater or even exclusive causal and therefore moral responsibility if their ultimate contribution to the injury is deemed more significant either because of the degree of traditional fault associated with this factor or because of what we perceive

as the greater metaphysical connection between that factor and its ultimate effect. That is, competing causal factors are assigned causal and therefore moral responsibility in part according to which factor has more traditional fault associated with it—the more unreasonable the act or omission happens to be the more causal and therefore moral responsibility we will assign to it. While fault can come in varying degrees, and the degree of fault can accordingly perform a helpful ranking function within the causal chain, an express finding of unreasonableness is not necessary.[45] A finding of traditional fault is not even a sufficient condition for finding causal and therefore moral responsibility, for even unreasonable conduct may not be assigned causal responsibility if it is too far removed from the injury at issue.

Suppose, for example, we decide to build a bridge. We know from past experience that a few workmen will inevitably be killed in the course of its construction, but we decide to build the bridge anyway. Let us hope that this decision was made only after weighing the benefits to be had from building the bridge against its cost in lives, but even if this was not the case, this does not mean that those who made the decision to build the bridge are necessarily to be assigned causal responsibility for the particular deaths that end up fulfilling our pre-existing statistical prophecy. When death strikes, there will most likely be other causal factors of much greater comparative significance further along the actual or counterfactual causal chain that led to any particular individual's demise (the safety harness provided to the worker who died failed, for example, or workers were not notified that gusts of high winds had been reported in the area, or they were allowed not to wear their safety

[45] Of course, I am not the first and I certainly won't be the last to argue that a traditional finding of fault—i.e. that the alleged wrongdoer either knew or should have known he was creating a risk of injury—is neither necessary nor sufficient for assigning moral responsibility. For a recent example of a such an argument, see George Sher, *Who Knew? Responsibility without Awareness* (Oxford: Oxford University Press, 2009), pp. 71-84.

harnesses if they did not want to, or the like). We are likely to treat the decision to build the bridge as simply part of what Hart and Honoré call "the stage already set," the taken-for-granted background upon which other factors act, and look to these factors as the most likely source of causal responsibility.[46] Given the closeness of these factors to the actual harm, they are better candidates for the assignment of causal (and therefore moral) responsibility because their adjustment is more likely to change the outcome in the future *in the relevant way;* that is, by allowing us to keep the benefits of building the bridge without the burdens. In other words, when assigning causal responsibility, mere causal contribution is not enough—those who merely authorized the construction of the bridge are most likely to be relieved of causal and therefore moral responsibility by more morally significant intervening causes, even if we consider their initial decision to build the bridge to be wasteful and frivolous. Not building the bridge would be like taking fewer at-bats; what we are really concerned about is whether there was a place along the chain of causation that most immediately led to these deaths where humans could have intervened and done something different to have prevented the deaths. And despite our ability to stipulate that no such opportunities existed in a hypothetical thought experiment, our real life experience suggests that in the real world we will always be able to find some.

Note that the inquiry that I have suggested should go on here is very similar to the inquiry we would undertake if we were trying to determine whether there has been an interference with someone's negative liberty. As Isaiah Berlin notes, when we are looking into the possibility that there has been such an interference, we are looking for

[46] Hart and Honoré, *Causation in the Law*, pp. 80, 179.

evidence of restraints on liberty generated by human agents.[47] To determine this, one needs a theory of causation, of what causes what, in order to determine whether the particular restraint at issue is directly or indirectly attributable to actual or counterfactual human action. Determining whether that restraint is attributable to human action, in turn, does not depend whether those potentially responsible for the restraint knew or should have known that their action would have such an effect—a finding of traditional fault is neither necessary nor sufficient for concluding that a restraint is the result of a human agent. Yet that does not make the conclusion that there has been such interference any less morally significant. If there is casual responsibility, there is necessarily moral responsibility too. Otherwise, there would be no need for such interferences to be *justified,* as negative liberty theorists plainly think there is.[48] Whether this moral responsibility triggers praise or blame or neither (in the case where the interference is permissible but not obligatory or prohibited) is then another matter, determined by other factors, but the issue of moral responsibility is settled by the causal inquiry itself.

This is also why it is not a response to my argument to point out that defenses only apply to moral responsibility and therefore moral responsibility and causal responsibility cannot both be (entirely) part of a single integrated inquiry. Defenses that are justifications go to the nature of what someone is responsible for, not to whether they are responsible. In other words, those considerations tell us whether what has been done is a moral wrong, and if so, how bad a moral wrong it is, and if not, whether it deserves moral praise and if so how much. They do not have any effect on our attribution of either causal or moral responsibility. If I shoot someone

[47] See Isaiah Berlin, "Two Concepts of Liberty," in *Liberty* (Oxford: Oxford University Press, 2002), pp. 169-170.

[48] See. e.g., Ronald Dworkin, *Taking Rights Seriously* (Cambridge: Harvard University Press, 1977), pp.266-278, esp. p. 267.

who is about to wrongfully shoot you, for example, I am both causally and morally responsible for the other shooter's death no less than I would be causally and morally responsible for your death if I shot you in a fit of unprovoked rage; it's just that in the former case I deserve thanks and praise rather than blame and punishment.[49]

Defenses that are excuses, in turn, go to *both* causal and moral responsibility, reducing the relative importance (but not the absolute existence) of each (or not, if they are unsuccessful) to the same degree. For example, duress is not a defense to murder, but it does show us that someone else is both causally and morally responsible for the victim's death, perhaps even more so than the person who carried out the killing. Similarly, if I break into your mountain cabin to shelter from an unexpected storm I am less casually and morally responsible for the damage to your cabin than if I had broken in for kicks on a sunny day; but I am still causally and morally responsible for the damage caused nonetheless; otherwise I would have no obligation to pay for the damage done in the former case and most people think that I do. Excuses reduce the all-things-considered importance of a particular factor's causal and therefore moral responsibility, but not the *pro tanto* existence of that factor's causal and moral responsibility; justifications change the nature of the act from one that is a *prima facie* wrong to one that is all-things-considered not, but in either case

[49] One reader of this paper suggested another possibility: that when I have acted in defense of another, I am morally responsible for the death of the potential assailant but not causally responsible for it (the assailant, presumably, being causally responsible in this case for his own death). But while the assailant may be causally (and indeed morally) responsible for his own death, the assignment of responsibility to the assailant does not reduce much less eliminate the responsibility of the Good Samaritan here for his righteous act. That is the whole point of my argument here—justifications change the character of the act that one is causally and morally responsible for; they do not reduce the degree of that causal and moral responsibility. Indeed, I do not see how the Good Samaritan could be morally but not causally responsible for the assailant's death, for it makes no sense to say someone is deserving of moral praise for something they are not causally responsible for. In any event, the issue I am addressing in this paper is not whether one can be morally responsible for something that one is not causally responsible for—for the reasons I have already given, I take it to be relatively uncontroversial that this could not be the case. The issue I am addressing in this paper is whether one can be causally responsible without being morally responsible. A lot of people assume this is possible and perhaps even common. It is this assumption that I am challenging in this paper.

causal and moral responsibility rise and fall together.[50] Indeed, whether a defense actually reduces causal and moral responsibility or merely determines the nature of our reactive attitude toward what it is that someone is responsible for may be a handy way of distinguishing a justification from an excuse.[51]

V. WHAT ABOUT OUGHT-IMPLIES-CAN?

One implication of my argument, of course, is that people can be held causally and therefore morally responsible for wrongs that it seems they had no reasonable opportunity to avoid. This brings us back to the original objection to my claim: that even if we actually do this, it cannot be morally right to do so. How can we dispense with a finding of traditional fault and still find moral responsibility? How can mere causal responsibility be enough to assign blame and impose punishment? How can holding a human factor causally responsible not only trigger the obligations of an insurer but also justify moral guilt and criticism and blame? Is this not unjust? In the absence of a finding of traditional moral fault, it seems we are imposing responsibility

[50] One reader of this paper suggested that the insanity defense may be an example of a defense where causal and moral responsibility come apart, for in cases where that defense is made out we would not assign moral responsibility to the insane person but we would still assign causal responsibility to him. After reflecting on this example, however, my view is that the causal responsibility we are referring to in this example is not the kind of causal responsibility I have been talking about throughout this paper. When someone is insane, just like when they are an infant or otherwise mentally incompetent, what we have decided is that they are not capable of bearing moral responsibility. In that sense they are like a special kind of thing, and when we assign causal responsibility to a thing we are doing something very different than when we are assigning causal responsibility to a being capable of bearing moral responsibility. We are not making a simple attribution of the metaphysical relation between two events; rather, we are making a far more complex determination about the relation between two people. So we might indeed assign causal responsibility to an insane person who kills another, but this is a different kind of causal responsibility than the causal responsibility we assign to those who are mentally competent. The insanity example therefore does not suggest that the kind of causal responsibility I have been talking about in this paper (the causal responsibility of a being that is capable of bearing moral responsibility) and moral responsibility come apart. Given the way I am using these concepts, they still rise and fall together.

[51] For further discussion of what distinguishes a justification from an excuse and citation to some of the vast literature devoted to this topic, see Vera Bergelson, "Justification or Excuse? Exploring the Meaning of Provocation," *Texas Tech Law Review* 42 (2009-2010): 307-319; Douglass Husak, "Partial Defenses," *Canadian Journal of Law and Jurisprudence* 11 (1998) 167-192. Note also that this same basis that I am proposing for distinguishing a justification from an excuse can also be used to explain the difference between blameworthiness and wrongfulness.

on someone who had no reasonable opportunity to conduct themselves in any other way.[52] It seems we have violated the principle that says ought implies can.

Unfortunately, a full discussion of the intricacies of the ought-implies-can principle would take a paper in itself, so I will not be able to give this issue all the attention that it deserves here. But I will point out that it is not clear that the ought-implies-can-principle even applies. We are not talking about *physical* or *logical* impossibility here, for if there is causal responsibility, then it must be both physically and logically possible for us to intervene in the causal chain in the proposed way. We are not even talking about *epistemic* impossibility here, for it will always have been possible to have discovered the information necessary for us to have taken preventive action, for that is what we have ultimately done by identifying the causal factor at issue as a candidate for causal (and therefore moral) responsibility. In other words, it will always have been possible to have known about the risk of injury to others that an activity posed, even if this was not actually known at the time, because we have by definition discovered this after the fact—otherwise the activity in question wouldn't be a candidate for being in the causal chain. So the only possible basis for concern here arises if we interpret the ought-implies-can principle broadly so as to require not merely epistemic *possibility* but epistemic *feasibility* too. If there was actual knowledge of this risk, there is a question as to whether the risk was reasonably encountered, and if there was not actual knowledge of the risk, there is a question as to whether such knowledge was reasonably discoverable at the time.

But if I am right about my assessment of human nature, about our unwillingness to accept that we live in a dangerous world, both of these questions are actually determined by the fact that serious injury is the result: any activity that ends

[52] See H. L. A. Hart, *The Concept of Law,* p. 173.

in serious injury is not reasonable in this sense, nor is our lack of knowledge or information. Given this aspect of our nature and the attitude toward the risks of life it promotes, it is simply *unreasonable by definition* for the causally responsible to causally contribute to a serious injury, no matter how many precautions the causally responsible might take. The ought-implies-can principle accordingly provides no bar to the argument I am making.[53] Given human nature and the conditions and circumstances under which we live, there is always causal responsibility to be placed somewhere, and if there is causal responsibility, there is always moral responsibility too.

Put another way, my point is to deny that reasonableness plays the deciding role in our moral reasoning that we think it does. It is a factor, and indeed a very important factor. But what I have argued is that the causally responsible will always be found to have acted unreasonably, and that sometimes one can be causally responsible even though we have not acted very unreasonably at all. In these cases, then, "unreasonableness" is not so much an *input* that goes into determining that we are morally responsible for an injury; rather it is an *output* of the decision that we are causally responsible and therefore morally responsible for that injury, and in some cases this determination is influenced primarily by factors other than the nature of our behavior. You can be causally responsible for a serious physical injury no matter how many precautions you might take, and if you are you will be morally responsible too, for in such cases there is no such thing as accidents.

[53] For a similar argument against the ought-implies-can principle, albeit one that proceeds by a slightly different route and reaches a slightly different but not inconsistent conclusion, see Matthew H. Kramer, "Moral Rights and the Limits of the Ought-Implies-Can Principle: Why Impeccable Precautions are No Excuse," *Inquiry* 48 (2005): 307-355, esp. 328-331 (note that Kramer argues that proof or fault is not required to establish a moral wrong; I argue that it is always present even when it is not technically required).