

The Consequences Matter-- But To What?

Barbara H. Fried, Stanford Law School (bfried@stanford.edu)

To: Participants in the Global Justice/Political Theory Workshop
From: Barbara Fried
Re: Session on December 11, 2009

I look forward to hearing your thoughts on the attached paper at the session on Friday.

The paper is part of a larger project on how we should regulate conduct that is socially productive, but poses some risk of harm to others. The official technique for risk regulation in the modern administrative state is some form of cost/benefit analysis: we tote up the expected social benefits and expected social costs of alternative courses of conduct, and opt for that course that is expected to generate the largest aggregate benefits (net of costs). There is a vast and growing critical literature on the normative, conceptual, and administrative problems with cost/benefit analysis. I agree with many of those criticisms. But my particular interest in this project is the objection raised to any aggregative procedure, on the familiar deontological ground that it fails to respect the distinct rights and interests of individuals. Accepting the moral impulse behind that objection as understandable, maybe even compelling, the question I want to address is whether nonconsequentialists have offered—or could offer-- any coherent alternative to aggregation in this context.

The short answer is, I don't think so. I want to underscore that this is not a broadside attack on nonconsequentialist principles or a categorical defense of consequentialism. It is a domain-specific concern. For reasons I touch on briefly in this chapter, I think that deontological principles by their nature are the wrong tools to solve the fundamental problems raised by accidental (unintended) harms, and that this reality has been obscured by the limited and somewhat peculiar focus of the immense philosophical literature on harm to others.

While these larger concerns are in the background, this chapter mostly focuses on the central deontological objection to aggregative solutions to regulating potentially harmful conduct: that it violates our duty not to act in a fashion that will result in harm to others.

The Consequences Matter-- But To What?

Barbara H. Fried, Stanford Law School (bfried@stanford.edu)

I. Introduction

In 1972, a new Ford Pinto stalled in the middle of a highway, causing a car behind it to rear-end it. The force of the collision pushed the Pinto's rear-mounted gas tank forward and into the differential housing, which punctured it. Fuel spraying from the punctured tank ignited inside the car, engulfing the car in flames. One of the two occupants died from the resulting burns. The other, a 13-year-old boy, was severely disfigured for life, requiring multiple skin grafts over the next decade.

In the lawsuit the heirs and survivors subsequently filed against Ford for defective product design, it emerged in pretrial discovery that before the Pinto ever went into production, Ford knew about the danger that a rear-end collision would explode its gas tank. Indeed, its own crash tests revealed the engine was vulnerable to puncture and explosion at crash speeds as low as 20 MPH. Internal memos from Ford engineers to top management estimated that at a cost of \$5 to \$11 per car, the Pinto could be redesigned to decrease the probability of explosions significantly, saving an expected 180 lives. Ford management declined to make the changes,

estimating that the changes would cost the company a total of \$137 million, but save the company only \$49.5 million in incremental damages for deaths, injuries, and property damage that could have been avoided through redesign.

The jury found against Ford, awarding the plaintiffs roughly \$3 million in compensatory damages, and the boy an additional \$125 million in punitive damages, an amount subsequently reduced on appeal.¹

In the years since, the Ford Pinto case has come to stand for all that is wrong with cost/benefit analysis-- the decision procedure that Ford allegedly employed in deciding not to redesign the Pinto. While the 'received' account of Ford's behavior continues to be disputed, if one takes that account to be more or less accurate one can easily understand the outrage it triggered.² Even on a straight cost/benefit analysis, Ford almost certainly made the wrong decision in not redesigning the car.³ The fact that the Pinto was significantly less safe than most other cars on the road was probably sufficient to make its behavior legally negligent even without a faulty cost/benefit analysis. The crash engineers at Ford strongly urged redesign, but were overruled by top management, who were notoriously indifferent to safety concerns. Finally, Ford was well aware that federal safety regulations then pending would require it to protect the integrity of the gas tank in crashes up to 30 MPH, but

¹ Grimshaw v. Ford Motor Co., 119 Cal.App.3d 757, 808 (1981), quoting Dawes v. Superior Court 111 Cal.App.3d 82, 88 (1980) hg. den.

² The best-known of these revisionist accounts is Gary T. Schwartz, "The Myth of the Ford Pinto Case," 43 Rutgers Law Review 1013 (1991).

³ See, e.g., <http://www.engineering.com/Library/ArticlesPage/tabid/85/articleType/ArticleView/articleId/166/Ford-Pinto.aspx>; Douglas Birsch, "Product Safety, Cost Benefit Analysis, and the Ford Pinto Case," in *The Ford Pinto Case: A Study in Applied Ethics* (Birsch and Fiedler, eds., 1994)

chose to wait until the regulations were adopted to comply with the standards. Add to all of that the fact that Ford withheld information about the risks for fear of bad publicity, and you have a poster child for bad corporate actors.

But the moral that many have drawn from the Ford Pinto case goes well beyond the particular facts of the case. Put simply, it is that trading off lives for any amount of money or other lesser 'goods' is wrong, because doing so fails to respect the essential worth of every human life. As one commentator put it, "It is difficult to understand how a price can be put on saving a human life.... [I]t seems unethical to determine that people should be allowed to die or be seriously injured because it would cost too much to prevent it."⁴ Any Kantian will recognize and endorse that sentiment, which echoes the central Kantian conviction that "human beings have dignity and not mere price," and the ideal Kantian legislator "unconditionally attribute[s] a worth to persons that cannot be quantified and is not subject to trade-offs."⁵ But so also will other nonconsequentialists, all of whom take us to have a duty to act in a fashion that will not harm others, and take that duty presumptively to trump pursuit of the general good.⁶

⁴ <http://www.wfu.edu/~palmitar/Law&Valuation/Papers/1999/Leggett-pinto.html>. Or as another commentator put it in more inflammatory prose: "Essentially, Ford argued before the government that it would be cheaper just to let their customers burn!"

<http://www.engineering.com/Library/ArticlesPage/tabid/85/articleType/ArticleView/articleId/166/Ford-Pinto.aspx>

⁵ Thomas E. Hill, Jr., "A Kantian Perspective on Moral Rules," in *Philosophical Perspectives*, 6, Ethics, 1992, at 292, 292

⁶ Hill, "A Kantian Perspective," at 292. "One cannot... justify disregarding or violating the dignity of a few persons with the thought that thereby one would promote more dignity in many others." *Id.* at 292. While rule utilitarians are always trying to "find the rules that produce the most utility, without prior constraint regarding its distribution or the means of achieving it," "Kantian legislators ... are trying to find rules they can reasonably endorse and justify to one another under the severe constraint of their overriding commitment to the dignity of each person." *Id.* at 298.

In the regulatory context, this view has led nonconsequentialists to reject not just cost/benefit analysis, but any aggregative procedure that trades off death or serious harm to one group against money or other non-vital benefits. But we make such tradeoffs all the time, in our private as well as collective decision-making capacities. Every time we drive a car, we put others' lives at risk, however cautiously we act. Every time we manufacture a product, however safe the design, for some amount of money we could make it even safer, and our failure to do so puts others' lives at risk. Is it really impermissible for us to impose such a risk of harm or death on others? If so, how would nonconsequentialists have us act instead?

It is indisputable that the harmful consequences of our acts matter, and that it goes "against the grain" -- against our universally shared intuitions-- to argue otherwise.⁷ As one commentator put it, "Harm matters. It has intrinsic significance."⁸ The question is, matters to what?

There are a number of possible answers to that question:

(a) Consequences matter psychologically and emotionally to all of us, without regard to any moral judgment about the action(s) that produced them or

For the balance of the paper, I will use the term "nonconsequentialism" to refer generally to all deontological principles (Kantianism, modern day contractualism, liberal rights theory) that take the individual's right to be free from serious harm to provide a Razian "exclusionary reason" that trumps any quantity of (qualitatively) lesser benefits to society. It is hard to draw clear distinctions among these theories for any purpose; for my purposes, the differences among them are indeterminate or immaterial.

⁷ Herman, id. at 95. "What I think our considered moral intuitions tell us here is simply that something untoward has happened which must have moral significance even if not ramifications." Id. at 97. One can find similar sentiments expressed throughout the nonconsequentialist literature. See, e.g., Anthony Duff, "Auctions, Lotteries and the Punishment of Attempts," *Law and Phil* 9 (1990): 1-37; Stone, "The Significance of Doing and Suffering"; Frances Kamm, *Intricate Ethics*, at 273; Julie Tannenbaum, "Emotional expressions of moral value," *Philosophical Studies* (2007), 132:43-57.

⁸ Ken Levy, at 290 n. 42.

the agents responsible for those actions. We live in the world of what actually happened, not what might have happened but didn't. Given that reality, it is hardly surprising that we would care deeply about the consequences of everyone's acts—our own and others'.

(b) Harmful consequences may matter to us in a humanitarian sense, insofar as they leave a victim in bad straits.

(c) Harmful consequences may matter to us in deciding what an actor owes the victim of her harmful conduct-- compensation, apology, efforts to repair, remorse-- *once we have concluded on other grounds that the conduct in question is wrongful (that is, criminal or negligent), or that, although the conduct is not 'wrongful' in this sense, that the actor should nonetheless be held strictly liable for the resulting harm.*

(d) Where liability is based on a claim of negligence, harmful consequences can be probative of whether the defendant used due care.

All of these ways in which consequences might matter figure in the nonconsequentialist literature, some of them prominently. Some seem unproblematic in principle, others less so. But I want to focus on a different, central sense in which the nonconsequentialist literature suggests they matter:

(e) The consequences of an act matter because they determine, *nunc pro tunc*, whether the conduct that produced them was permissible.⁹ That *ex post*

⁹ Here and throughout the paper, "conduct" or "act" refers to a chosen activity undertaken *with a specified level of precaution against harm to others*. Thus, "driving at 50 MPH" is one act; "driving at 65 MPH" is a different one. When I speak of the state "prohibiting" or "permitting" an act, I therefore mean prohibiting or permitting that activity when undertaken with the specified level of precaution.

judgment is translated into an ex ante command in the frequently encountered injunction, “Do not cause harm.”

Words like “permissible,” “wrongful,” “justified,” “excused,” have caused no end of conceptual trouble in this area, because they can mean so many different things. So I want to be clear on what I mean by “impermissible” here. I mean that we believe the actor, all things taken together, was morally required to have acted otherwise. To put this in legal rather than moral terms, we retroactively conclude that the conduct is *malum prohibitum*, meaning that for the good of society we wish it never to occur, and if we could costlessly enjoin it, we would do so.¹⁰

For the balance of the article, unless otherwise stated I will use “impermissible (or wrongful) conduct” and “conduct we wish to prohibit” interchangeably in this sense.

It is not always easy to tell in which of these senses (or others) nonconsequentialists believe that consequences matter. But there is no shortage of statements to suggest that they matter to nonconsequentialists in this last sense-- that an act can become impermissible solely in virtue of its bad consequences. To quote one famous instance, from Bernard Williams’ sometime defense of the morality of moral luck: If one takes a gamble and loses, “one did the wrong thing, in the sense that [one] has no basis for the thought that *[one] was justified in acting*

¹⁰ The mere fact that the state will not enjoin the conduct ex ante does not mean it is permissible, in the sense in which I am using the word. For a variety of prudential and principled reasons, we put a very high burden of proof on those seeking to enjoin another’s conduct, and in some areas—notably the 1st Amendment—almost never grant prior restraint orders. The qualifier “costlessly” is meant to capture these prudential and principled reasons not to enjoin conduct ex ante.

as [one] did."¹¹ What this could possibly mean, and what its implications are for public policy, are the central questions I want to address here.

I wish to make five points. The first four are developed in more detail in Parts II and III below, and the fifth I put on the table here as a placeholder for future discussion.

First, when we are contemplating acts, the consequences of which are uncertain *ex ante*, this view of why consequences matter requires us to suspend ultimate judgment about the permissibility of an act until it has been completed and all its consequences known. Whatever else this is-- and various possibilities are discussed below-- it is not an action-guiding norm. (By "action" here, I mean the initial act that caused the untoward consequences, not any subsequent actions one might take to repair or compensate for the harm.) There are two ways to interpret the injunction "Do not cause harm" that would convert it to a principle that in theory can be acted on: "Do not do anything that could possibly cause harm" or "Do not do anything if and only if it is certain to cause harm." But the first would rule out virtually all action, and the second virtually none.

I am hardly the first to make these observations, but their significance has yet to register adequately in the nonconsequentialist literature on harm to others. The principal reason for this, I believe, is that the canonical hypotheticals that have framed discussion in the literature on permissible harm to others have pushed the problem of uncertainty to the margins. Some have done so by viewing the problem of harm to others from an *ex ante* perspective—before the actor has committed to a

¹¹ Williams, "Moral Luck," at 23 (itals. added).

contemplated course of action—but stipulating the consequences that *will* (with a 100 percent certainty) ensue from a given choice. In other cases, while the conduct in question poses uncertain risks *ex ante*, the reader is invited to judge it from an *ex post* perspective, by which time the consequences are known with certainty. The effect in both cases is to suppress the difficulties that immediately arise if we try to scale up the moral intuition “Do not cause harm” to deal with the sorts of decisions we face every day-- garden-variety decisions that have upside and downside risks.

The “harm to others” literature is hardly alone in marginalizing the problem of uncertainty. As one commentator has noted, “Throughout the history of moral philosophy, moral theorizing has for the most part referred to a deterministic world in which the morally relevant properties of human actions are both well-determined and knowable.”¹² In the harm-to-others literature, that focus has produced a literature in which the tail (trolley-type problems and other one-off rescue cases) is wagging the dog (regulation of socially productive conduct that poses some risks of harm).

Second, when nonconsequentialists say we have a “duty not to harm others” (that is, a duty not to cause bad consequences to befall others) I believe that most of the time they really mean something else: that we have a duty to compensate the victim for the harm we did, apologize for it, feel regret about it, or to use due care in avoiding it. These are very different things from saying that we have a duty not to act in a way that (will?) (could?) result in harm to others. Most of the moral luck literature has been clear on this distinction, focusing on the apparent paradox of

¹² Hansson, “Risk and Ethics,” p. 30.

judging an act to have been morally deficient in virtue of its consequences while simultaneously not wishing the actor, “all things taken together, to have acted otherwise.”¹³ Not so the philosophical literature on permissible “harm to others.” There, confusion on this point has persisted in part because of the ambiguity of the terms that anchor this debate (wrongful, permissible, justified, etc.), and in part (I believe) because the ex post perspective from which the nonconsequentialist literature has generally viewed the problem of risk. From that perspective, talk about prohibiting risky conduct is cheap, in the sense that it is too late to prohibit the conduct in question whether we wish to or not; that horse has left the barn. Absent the need to specify what risky conduct is ruled out by the “duty not to harm,” it is easy to believe there must be some answer that does not rule out all action.

Third, many of the hypotheticals offered in support of the proposition that “the consequences matter” in sense (e) introduce confounding features that provide independent reasons to condemn the conduct in question-- independent, that is, of the mere fact that it produced bad consequences. These include the presence of identifiable victims, certainty about the incidence and extent of the harm, infliction of harm through the deliberate active agency of the injurer, infliction of harm as the “upstream” means to aid others, and conduct that everyone would agree is undesirable (negligent, criminal) irrespective of the consequences. The last of these aside, why any of these features should be morally relevant to the acceptability of an act has generated a huge literature I hope not to engage directly here. For these purposes, I will assume they are morally relevant. My concern is

¹³ Williams, *Moral Luck*, at 31.

different and simpler: that the presence of so many other factors that independently trigger the intuition that a given act is impermissible makes it hard to isolate how much work the consequences, standing alone, are doing in generating that intuition.

To get a handle on that question, we need to strip out those confounding factors. What we are left with is with the garden-variety problem of risk regulation under uncertainty: Should we permit conduct that has expected social benefits but also expected (if uncertain) costs in the form of future harm to as yet unidentified individuals. Should we permit a newly developed AIDS vaccine to be distributed, where clinical trials suggest it is likely to save 100,000 lives a year but result in one death from an adverse reaction? Should we raise the speed limit on the highway from 55 MPH to 65 MPH, where the change is projected to generate tens of billions of dollars a year in increased productivity from time saved driving, but result in 5,000 more accidental deaths? Should we let people fly private planes, knowing that 50 bystanders are likely to die each year from crashes? If so, what sort of training should we require before we license private pilots? What sort of safety features on the planes?¹⁴

These sorts of decisions—whether to permit socially productive conduct that poses some uncertain and distant risks of harm to unidentifiable others and what level of safety precautions to insist upon-- are not outliers. Outside of the criminal

¹⁴ It is impossible to eliminate the act/omission distinction,, as any choice can be argued to be one rather than the other, and the resulting characterization argued to have moral significance. But, as I will suggest below, whether authorizing the vaccine or driving at 65 MPH or flying a private plane with specified safety precautions is characterized as an act or an omission is unlikely to affect nonconsequentialists' handling of these cases. (See *infra*, pp.).

context, the vast majority of potentially harmful conduct regulated by the state takes this form. It is the trolley problems, Scanlon's technician trapped in a broadcasting studio, Taurek's choice between saving one person and saving five, Williams' Jim and the Indians, that are the freaks and sports of human social interaction. Thus, how nonconsequentialist principles would have us make such decisions is of more than academic (in both senses of the word) interest.

Fourth, in my view, the nonconsequentialist literature to date has not offered an analytically coherent *substantive* answer to this question that doesn't boil down to some form of aggregation indistinct from consequentialist solutions.¹⁵

The standard answer an act consequentialist would give is straightforward, in principle if not in practice: We should tote up the expected benefits and expected costs of not approving the vaccine v. approving it with various levels of safety precautions (increasing the number of clinical trials required, lowering the level of acceptable risks, etc.) and opt for the course that we determine by some aggregation procedure to have the highest expected social value. The nonconsequentialist objects to such a decision rule, *inter alia* on the ground that unless those who will

¹⁵ I set to the side procedural solutions, which offer a clear alternative to substantive decision rules--consequentialist and nonconsequentialist alike. See, e.g., Henry Richardson, "The Stupidity of the Cost Benefit Standard" in Adler and Posner, p. 163; Richardson, "Beyond Good and Right: Toward a Constructive Ethical Pragmatism," *Philosophy & Public Affairs* 24 (1995): 108-141. I also set to the side the very strong possibility that cost/benefit calculus or any other form of consequentialism cannot be operationalized without smuggling in policy-makers' own value judgments about what is worth optimizing. For a version of this concern, see, e.g., Viscusi and Hamilton; Viscusi, "Risk Equity" (on the Superfund project).

I use "aggregative procedures" throughout the paper to refer to all decision rules that evaluate conduct by balancing aggregate expected benefits against aggregate expected costs in some fashion, without giving any particular cost the force of an exclusionary reason. I do not mean to limit the method of aggregation or the "benefits" and "costs" that count to those employed in conventional cost/benefit analysis.

(are expected to?) bear the costs will (are expected to?) also enjoy a commensurate share of the benefits, such a rule impermissibly trades off one person's ex post well-being for another's.¹⁶ The question is: what substantive decision rule would they offer in its place to avoid this result?

Some nonconsequentialists, holding firm on the view that the duty not to harm others operates as a side constraint on aggregative solutions, have concluded that any course of conduct that will (certainly? predictably? possibly?) turn out to have harmed others is impermissible. That view, whatever other virtues it might have, cannot provide any guidance in deciding how to act oneself, or what conduct we should collectively permit others to engage in. Other nonconsequentialists have qualified that absolute duty not to harm with the additional requirement that the action must also be "wrongful," "negligent," unjustified," "without excuse," etc. Where those qualifications are given any operational meaning, it is hard to distinguish from some form of cost/benefit calculus. Still others have put off the permissibility question for another day, focusing their energies instead on defending a broad obligation to compensate for any harm caused by actions adjudged (on as yet unspecified grounds) to be impermissible.

Finally, a sizeable group has explicitly ceded the domain of decisionmaking under uncertainty to the consequentialists, to resolve in accordance with some version of a cost/benefit calculus. But, as my comments above suggest, this is no minor concession. Once we set aside criminal conduct and intentional torts

¹⁶ Of course, nonconsequentialists object to cost/benefit analysis as customarily practiced on numerous other grounds besides aggregation. Many of these objections are telling, if not dispositive, but they are not the subject of analysis here.

(conduct that generally everyone can agree is impermissible, even if they can't agree on the reasons), virtually all conduct poses some expected social benefits and some uncertain risk of harm, if judged ex ante—the only epistemological perspective from which we can decide how to act. This is true whether legislatures or administrative agencies step in ex ante to prohibit certain kinds of conduct (driving over 55 MPH, distributing a new drug without getting FDA approval, etc.), courts step in ex post to define what constitutes negligent behavior through a common-law decision making process, or individuals are trying to decide for themselves how to act.

If the view that “consequences matter” has no guidance to offer legislatures, administrative agencies, courts or individuals engaged in these tasks, it is of limited policy relevance. Were that failure attributable to the gap between ideal and non-ideal theory, nonconsequentialist principles need not be adjudged a failure for attending only to the former; there is a legitimate role for division of intellectual labor here. But that is not the source of the problem. Ideal theory as much as nonideal theory concerns itself with how we should act in this world-- not just what we owe to others whose rights we have trampled, whose personhood we have failed to respect, etc., but how to act so as to avoid trampling those rights or violating that personhood in the first place. If nonconsequentialist principles have nothing to say on the latter question for the vast majority of conduct in this world, they are not just of limited policy relevance in the domain of harm to others; they are of limited moral relevance as well-- unless, of course, we are to treat the ubiquity of risk of harm as itself a nonideal feature of the world. That possibility brings me to the last point.

Fifth, I believe there are reasons to be skeptical that nonconsequentialist principles *can* generate a substantive decision rule for determining the permissibility of socially useful but risky conduct that avoids trading off the individual “right not to be harmed” for other social goods. Fleshing this point out and doing justice to the counterarguments is too large a project to undertake here, and I will not pursue the point further in Parts II or III. But I want at least to put the reasons for my skepticism on the table, in the hope they might provoke further discussion.

Read as an aspirational norm, the “duty not to harm others,” like the broader nonconsequentialist principles from which it is derived (the Kantian kingdom of ends or universal law, the classic liberal structure of rights, what people would reasonably consent to under some versions of contractualism) presupposes a natural and social world in which we can generally advance our own *prima facie* legitimate interests without compromising others’. It imagines a world in which, in Locke’s words, people are “naturally in... a state of perfect freedom to order their actions, and dispose of their possessions and persons, as they think fit, ... without asking leave, or depending upon the will of any other man,” subject only to the requirement that, all of us “being... equal and independent, no one ought to harm another in his life, health, liberty, or possessions.”¹⁷ Only in such a world could we

¹⁷ Locke, Second Treatise, secs 4 and 6.

value others' "rational capacity *unconditionally*," or act so as to "cause *no* untoward consequences" to others.¹⁸

This optimistic view of human coexistence bears limited relationship to the social world we live in. In the real world, tradeoffs among individuals' legitimate interests are ubiquitous. (Were they not, there would be no need for the government to regulate conduct to begin with.) They are just invisible to most people until something goes wrong. You have to decide whether to drive or walk to the store, which is a mile from your home. You know that each time you get behind the wheel, there is a very small chance that you will end up hitting a pedestrian, no matter how careful you are. The mile walk is relatively inconvenient to you; it is also raining hard, which makes walking unpleasant (but also slightly increases the risk of a car accident if you drive, even very carefully). Do you have a right to drive? Either way we decide it, we are, prospectively, trading off one person's interests against others'.

Such tradeoffs are not limited to potential physical harm to others. They are endemic to almost every part of social life, starting with the rule of law itself. All rules are necessarily over- and under-inclusive with respect to the ends they seek to advance. When we set the minimum driving age at 17, we are not only trading off the interests of young drivers against those they might potentially harm by virtue of their immaturity; we are also trading off the interests of different groups of young drivers. Some teenagers have the judgment to be safe drivers at age 14, and have

¹⁸ Julie Tannenbaum, "Emotional expressions of moral value," *Philosophical Studies* 132:43-57, 48 (2007) (itals added); Honore, "The Morality of Tort Law—Questions and Answers," in Owen, ed., *Philosophical Foundations of Tort Law* (Clarendon Press, 1995), 73-95 at 75 (itals added)..

their autonomy severely constrained by being made to wait until they are 17; others will be serious dangers to themselves and others until they are 21 (or beyond). Because we lack the information that would let us distinguish one group from another, we must subject them all to a uniform rule, thereby sacrificing the tails of the distribution (the legitimate autonomy interests of the mature 14-year-olds and the safety of the victims of those who will not be responsible drivers until they are 21) to the interests of all others.

These ubiquitous tradeoffs are the tragedy of social life; the question is, do the essentially optimistic premises of nonconsequentialism have the resources to cope with it? I am doubtful.

It is tempting to think we can avoid interpersonal aggregation by viewing the collective decision whether to permit a given act from some variant on Rawls's original position, which in effect forces each individual to internalize the costs and benefits of every course of action, thereby aligning each individual's interests with the optimal aggregative solution. But that move proves too much and too little, for reasons familiar from the social contractarian literature.¹⁹ If we posit complete ignorance about one's ex post position and a psychologically plausible degree of risk aversion, we end up with something like Harsanyi's solution, adjusted to allow for individual differences in preferences. Each person will choose the solution that

¹⁹ This comment pertains only to the decision whether a given risky act is permissible to undertake, not to the question whether anyone ultimately harmed by the conduct might have a right to compensation from individuals or from the state. Fairness concerns, whether posed from a contractualist perspective (e.g., George Fletcher's argument) or some other nonconsequentialist perspective (e.g., the arguments of Jules Coleman, John Goldberg, Benjamin Zipursky, Ernest Weinrib et al. in favor of a corrective justice approach to tort law) may well have a lot to say on the latter issue.

optimizes her expected value, the values she attaches to different outcomes and her degree of risk aversion.²⁰ Assuming a plausible range of risk aversion, almost everyone will be deemed to have assumed the risk of a wide range of accidental harms. Alternatively, if we posit extreme risk aversion of a Rawlsian sort or—what amounts to the same thing— give those who, ex post, will turn out to be gravely harmed prospective knowledge of that fact and a veto over the conduct in question (Scanlon’s sometime solution), then we end up with the answer that virtually no conduct that poses a risk of grave harm to at least one person is permissible, provided there is some alternative course of action that could avoid grave harm to anyone.²¹ For how many of us would volunteer for certain death so that their fellow citizens might avoid chicken pox, have a new bridge to ease traffic congestion, not have to lose twenty minutes of a soccer game broadcast, or even avoid death themselves?

²⁰ See, e.g., Judith Thomson, *The Realm of Rights*, ch. 7; McCarthy, “Actions, Beliefs,” at 64-65; Sophia Reibetanz, “Contractualism and Aggregation,” at 301-02. The Rawlsian veil of ignorance would deprive individuals of any particularized knowledge of their expected ex post position, whether that knowledge would normally be available to them in the real world or not. Scanlonians, committed to creating thicker selves who are parties to the hypothetical agreement, have sometimes opted for what Reibetanz has called a “natural veil of ignorance,” meaning we do not “take from complainants knowledge which they would otherwise have,” but do not endow them with knowledge they otherwise would *not* have, about “whether or how general acceptance of the relevant principles [will] affect them.” Reibetanz, “Contractualism and Aggregation” at 301 n. 9. This is one place we could split the difference; there are many others. For my purposes, the critical point is that, whatever the basis on which the would-be complainant is endowed with or deprived of knowledge that her ex post position will be considerably worse than any one else’s, it is the presence or absence of that knowledge that will determine whether she will be deemed to reasonably reject the aggregatively optimal solution or not.

²¹ See, e.g., Scanlon, *What We Owe to Each Other*, p. 207: “In assessing the rejectability of [a] principle” based on burdensomeness, “we can begin...by taking the maximum level of burdensomeness and asking whether that would give a potential agent reason to reject the principle.” Applying that procedure to a standard risk regulation case, Scanlon concludes that where the activity in question has benefits for many people but “involve[s] risk of serious harm to some.... [i]t is obvious what the generic reasons would be for rejecting such a principle from the standpoint of someone *who is seriously injured*.” *Id.* at (itals added). See also Reibetanz, “Contractualism and Aggregation” at 302-04.

Much of the nonconsequentialist literature alternates between these two extreme positions, opting for the ex ante perspective when explaining why a particular risky act is permissible and the ex post perspective when explaining why it is not, often without supplying a meta-principle that explains when we are entitled to help ourselves to one perspective or the other.

II. Is it really the consequences that matter?

As suggested above, the canonical hypotheticals in the harm-to-others literature typically contain a number of features that independently trigger the intuition that a particular act is impermissible—independent, that is, of its consequences. The question is, if we strip out all these other factors, how much do the consequences simpliciter—“the fact that one has harmed the other”—matter to people?²²

1. Certainty about the incidence and nature of the harm. As noted above, the nonconsequentialist literature has generally approached the problem of harm to others from perspectives that remove or suppress uncertainty about the consequences of alternative courses of action.

In some cases, we are asked to choose a morally acceptable course of conduct from an ex ante epistemological perspective, but the outcomes of the relevant choices are stipulated to be certain. The standard rescue hypotheticals (the trolley problem oeuvre, Scanlon’s choice between saving a technician from bodily harm and interrupting the viewing enjoyment of millions of soccer fans, Taurek’s choice

²² Kagan, *Causation and Responsibility*” at 294.

between saving one and saving five, Williams' Jim and the Indians) all fit this mold.²³

If the consequences of an act are known with certainty *ex ante*, then expected consequences equal actual consequences, and the command to take actual consequences into account in judging the permissibility of an act, whatever else it is, is not paradoxical.

In other cases, we are asked to assess the permissibility of an act from an *ex post* epistemological perspective, when its bad consequences are known. The tort system generally approaches harmful conduct from this posture.²⁴ Many hypotheticals in the philosophical literature do as well-- indeed, some nonconsequentialists have gone so far as to argue that we *cannot* assess the permissibility of conduct until we know its consequences (more on this point below).²⁵ Of course, in almost all of these cases, the consequences of the same act, judged *ex ante*, were uncertain. One inadvertent effect of adopting the *ex post* perspective, I believe, has been to suppress that reality by conflating the *ex ante* and *ex post* epistemological perspectives.

²³ Thomson; Taurek; Scanlon; Williams. For other examples, see Sophia Reibetanz, "Contractualism and Aggregation" at 302-05 (the exploded mine hypothetical); Rahul Kumar, "Defending the Moral Moderate: Contractualism and Common Sense," *Phil. And Pub. Affairs*, 294-95 ("one person (call her Allie) will lose her life unless another person (call him Geoff) loses his arm...What reasons can we offer her in defense of the validity of a principle compliance with which she knows will result in her death, while compliance with some other principle will result in her living and someone else's merely losing his arm?")

²⁴ Contrary to frequent claims in the philosophical literature, potential victims generally have a right to seek an *ex ante* injunction to prohibit harmful conduct from occurring in the first place, and in some areas of the law such injunctions are routinely sought and granted. See, e.g., . The fact that most tort suits are filed *ex post*, after the harm has already occurred, results not from any principled bar on *ex ante* injunctions, but simply from the fact that few potential victims are able to meet the standing requirements to obtain an *ex ante* injunction (showing that injury to the plaintiff in particular is very likely to occur, that the injury will be substantial, and that it cannot adequately be redressed by *ex post* compensation), or indeed are even aware that they are threatened with harm.

²⁵ See discussion at TAN

Consider the following. In “Moral Luck,” Williams asks us to imagine the agent regret experienced by “the lorry driver, who, through no fault of his, runs over a child.”²⁶ One commentator, examining the deficiencies of a “deliberative-Kantian view” in handling a hypothetical like Williams’s, restates the hypothetical as follows:

[I]f an agent unknowingly kills another without making efforts to determine whether his action *would* have this result, his action is morally wrong[, because] the agent’s deliberations ... do not manifest that the agent values rational capacity unconditionally....

[But the lorry driver] pursues his job by driving down this street on the condition that and with the belief that his action *will not* kill anyone. Moreover, he has and pursues the sub-end of being aware of whether his action of driving *will kill anyone*. The trouble is that he arrives at the *false belief* that driving down this street *will not kill anyone*. His belief is false but reasonable— he meets all the reasonable requirements for justification in forming this belief.... [From a deliberative-Kantian view] then the driver’s action of killing the child is not morally deficient, since the action results from a decision that reflects that the agent values rational capacity unconditionally. ²⁷

Note what has happened in the elaboration of Williams’s hypothetical: The lorry driver’s *ex ante* belief about the outcome has been stated as a certitude (“he *will not* kill anyone”). This belief is then described as *false* in virtue of the fact that *ex post* he does in fact kill someone, but nonetheless reasonable in the sense that he did what he reasonably could to ascertain the *correct* answer ahead of time.

But surely this formulation misstates the situation in crucial respects. If asked *ex ante* to articulate his beliefs, the verbally punctilious lorry driver would say, “Because I plan to drive prudently, I believe I am very unlikely to kill anyone.” Stated that way, his belief is not falsified in any respect by the fact that *ex post* he

²⁶ Moral Luck at 28.

²⁷ Julie Tannenbaum, “Emotional expressions of moral value,” at 48.

does kill someone, and he succeeds in achieving his sub-end-- to drive in a fashion that will minimize the likelihood of running someone over—notwithstanding that he kills someone. This reformulation raises the question whether the entire “problem” of explaining agent-regret for bad outcomes in Kantian terms is a non-problem, born of mis-specifying what the “sub-end” is (and ought to be) for actors like the lorry driver, acting under conditions of uncertainty. But for present purposes, I want to make a simpler point: By recasting what is really a decision made under conditions of uncertainty in terms that eliminate the uncertainty, the author inadvertently invites the reader to conclude that what actually happened was inevitable, or at least much more likely, *ex ante*, to come to pass than it really was, and that the driver somehow could and should have done more than he did to avoid it, notwithstanding that he is stipulated to have acted reasonably. This is, of course, the problem of hindsight bias.

The one scenario that has not figured prominently in the nonconsequentialist literature is conduct, the consequences of which are uncertain *ex ante* and the permissibility of which we are asked to judge from an *ex ante* epistemological perspective. One would expect the contractualist literature-- committed as it is to an *ex ante* perspective stripped of any particularized knowledge of outcomes-- to be an exception. Sometimes it is, and when it is the consequence is generally to push the author to the aggregative solution.²⁸ More often, however, contractualists have induced some of the same psychological sense of *ex ante* certainty about the

²⁸ See, e.g., Scanlon, ; Reibetanz, “Contractualism and Aggregation” at 301-02.

consequences by focusing our moral attention on the hypothetical person who will be worst off ex post.²⁹

In addition to increasing the psychological salience of harm and inviting hindsight bias about its probability, the nonconsequentialist literature's focus on conduct with known consequences has tapped into people's well-documented propensity—psychological or moral—to weigh harms more seriously when they befall known victims, or when we believe them to be absolutely certain, rather than possible, probable, or even statistically certain, to occur.³⁰

2. The victim is identifiable, and the actor inflicts harm actively, knowingly, 'up close and personal' and/or as a means of benefitting others.

Virtually all of the standard hypotheticals in the nonconsequentialist literature on harm to others contain a particularized (as opposed to statistical)

²⁹ Rawls does this indirectly via the assumption of extreme risk aversion, which in turn invites us to imagine ourselves as the person who will be least well off ex post. Scanlon does it directly, by requiring that we (hypothetically) justify our action to the fictive person who, in the event, will have the most serious complaints ex post as a consequence of that action. For further discussion, see pp. infra.

³⁰ For arguments that implicitly treat as morally significant the difference between an outcome with an expected probability of .99 and one with a probability of 1.0, see Kamm, *Intricate Ethics* at 273; Scanlon, "Contractualism and Utilitarianism"; Sophia Reibetanz, "Contractualism and Aggregation" at 303-04 (measuring a complaint by the expected value to each person of adopting a principle whenever "we do not know [ex ante] whether anyone will be affected in a certain way by acceptance of some principle," but measuring it "based upon the full magnitude of that harm or benefit: whenever "we know that someone will be harmed or benefited in a certain way," even if we can't "identify the bearer of this complaint in advance.")

The prevalence and seductive force of an ex post epistemological perspective accounts, I believe, for another oddity in the nonconsequentialist literature: the widespread belief that harming someone is a different act from doing something that imposes a risk of harm, rather than the same act judged from different temporal perspectives. Confusion on this point has, among other things, led many nonconsequentialists to believe that they can cede the regulation of risky conduct to welfarists while preserving a significant domain of harmful conduct to be resolved by nonconsequentialist principles. See infra at p. ; B. Fried and M. Kelman, "Risk and Rights" (draft ms.).

victim. They also require the would-be harmer to inflict the harm through deliberate acts that are undertaken with full foresight of their consequences, to witness the resulting harm ‘up close and personal.’ In some, the harm must be imposed as a necessary means to benefit others, rather than an unavoidable collateral consequence of actions that benefit others (“upstream” versus “downstream” harms).³¹

Whether the salience of these factors is best explained in psychological or moral terms, and if the latter, why they are morally relevant, have been debated extensively in the philosophical literature. For present purposes, I will assume that these factors deserves some moral weight, even if not the full weight they have sometimes been given. Instead, I want to make two points.

First, the ubiquity of these confounding factors in the hypotheticals that dominate the philosophical literature makes it difficult to isolate how much the *consequences themselves* matter to the permissibility of conduct. Consider in this regard Scanlon’s hypothetical about medical experiments. Would it be all right, Scanlon asks, for a community to choose by lottery a small number of its own members to be subjects of medical experiments, where the aggregate benefits from those experiments will significantly outweigh the burdens on the unlucky few chosen to endure them?³² Scanlon expresses doubt, “because of the severe burdens it involves” for those unlucky few. But surely the burdens alone are the least of

³¹ Almost all nonconsequentialists balk at the imposition of upstream harms. See, e.g., Kamm, *Intricate Ethics*; Kumar, “Defending the Moral Moderate,” at 304-05 (“A moral system that allows one person to consider using another person... in the same way as one would use a convenient stick of wood to bludgeon the oncoming dog...amounts to...treating another as an object, not a person.”)

³² *What We Owe to Each Other* at 208-09.

what would bother most people here. Other factors likely contributing to Scanlon's intuition-- and, one suspects, without which the intuition would disappear—are the fact that the harm will occur as a means to an agent's end rather than a downstream consequence; the victims will be identified prior to the harm-generating act, thereby forcing the community to inflict pain on known victims as well as creating the opportunity for the victims to renege on their agreement, and the harm will be intentionally inflicted by active human agency, up close and personal.³³

Second, whatever moral weight these and other “how” factors deserve in determining what harmful consequences we may visit on others, they are irrelevant to the garden-variety decisions we face every day, as individuals and as a society, that require us to trade off one group's wellbeing for another's. (Should the speed limit be 55 or 65 MPH? Should we market the AIDS vaccine? Should I drive even though it is raining heavily?). As in the case of the construction project, we are typically trading off risk of future harm to as yet unidentified, statistical victims against benefits for others, often themselves unidentified. As in the case of the construction project, the act/omission distinction either doesn't map on to the available choices (are we “failing to prevent harm” or “knowingly inflicting it” when we raise the speed limit from 55 MPH to 65 MPH?) or produces results most people would find unacceptable (we cannot build a water treatment plant, because the process of building it is expected to *kill* one innocent person, while not building it

³³ Scanlon himself acknowledges the force of these confounding factors when he turns to the hypothetical of Jones, trapped under equipment in a TV transmission facility. See *infra*, at .

merely fails to prevent the deaths of thousands of others from contaminated water.)³⁴ Moreover, where the state rather than an individual is the relevant actor, the distinction between act and omission arguably loses whatever moral force it might otherwise have.³⁵

3. The conduct is negligent or otherwise undesirable irrespective of the consequences. Finally, many hypotheticals involve conduct that (in legal terms) is *malum in se* or reckless.³⁶ Other cases state or strongly imply that the harm in question was easily avoidable, the consequence of simple carelessness—in legal terms, negligent.³⁷ In all such cases, virtually everyone would agree the conduct is impermissible, irrespective of the consequences. Since nothing turns on the difference, it is easy to conflate a “duty to use due care not to cause harm” with a “duty not to harm.”

³⁴ In the case of vaccination policies, the preference for failing to aid over actively harming has often led to inaction (don’t release vaccines or mandate vaccination) where action would have produced much better outcomes. See Baron and Ritov, 1993.

³⁵ For a forceful argument to that effect, see Thomas Nagel, *Equality and Partiality*, 84 (1995).

³⁶ Interestingly, the legal academy has also given a central role to intentional harms (which really belong with criminal rather than tortious conduct) in its taxonomy of torts. The result has been to marginalize the core cases handled by the tort system: cases of accidental harm. For an excellent historical account of this, see Thomas Grey, “Accidental Torts,” 54 *Vanderbilt L Rev* (2001).

³⁷ The language of negligence repeatedly sneaks into accounts of our “duty not to harm.” See, e.g., Herman, “What Happens to the Consequences,” at 111 (“I am morally accountable for knowing that what I propose to do ... will injure someone or will involve breaking a promise—at least when the *wrong or injury is in plain sight.*” (itals added)). Goldberg and Zipursky, “Unrealized Torts,” at 1697 (one has a moral duty not to cause “physical injury [through one’s] *misfeasance.*”). For further discussion on this point, see Fried, “Is There a Coherent Alternative to Cost/Benefit Analysis?” (draft ms. 2009). As discussed below, discussions of the duty to compensate for harms one has caused suffers from the same confusion between causation and fault.

The question is, would nonconsequentialists' intuitions about the permissibility of harm-producing conduct change significantly if we stripped out all of these confounding factors? I believe so. By way of illustration, consider more closely Taurek's and Scanlon's famous hypotheticals rejecting aggregative solutions:

(i) Taurek: Six people are stranded on rocks in the ocean, five on one and one on the other. The tide is rising, and all will certainly die unless they are rescued very soon. You happen upon them, but have only enough time to get to one of the two rocks before the tide will drown them all. What should you do?³⁸

Taurek, famously, argued it would be wrong to save the five based on their greater numbers, because to do so would fail to show equal concern and respect for the one. The only decision rule that would "express [our] equal concern and respect for each person," Taurek argued, is to flip a coin as between saving the one and saving the five, which would give each of the six a 50% chance of survival.³⁹ Contra Taurek, almost all nonconsequentialists would save the five, but like Taurek, none justifies the preferred solution by a simple appeal to the numbers.⁴⁰

(ii) Scanlon: Jones is working in a broadcasting room of a TV station. Electrical equipment falls on his arm; the only way to rescue him is to shut off the

³⁸ John Taurek, "Should the Numbers Count?" *Philosophy and Public Affairs* 6: 293-316 (1977). See also Michael Otsuka, "Scanlon and the Claims of the Many Versus the One." *Analysis* 60(3): 288-93 (2000).

³⁹ *Id.* at 303.

⁴⁰ Some have justified saving the five by interpreting the principle of "equal respect" to require rather than forbid saving the five. Others have argued that the numbers should count, but only as a tiebreaker where the only thing that differs between the two groups is the numbers. See, e.g., F. M. Kamm, *Morality, Mortality, Volume I: Death and Whom to Save From It* (Oxford University Press, 1993), chs. 5 and 6; Scanlon, *What We Owe to Each Other*, ; Derek Parfit, "Innumerate Ethics." *Philosophy and Public Affairs* 7(4): 285-301 (1978); Gregory Kavka, "The Numbers Should Count." *Philosophical Studies* 36: 285-294 (1979); John T. Sanders, *Why the Numbers Should Sometimes Count*, *Phil. & Pub. Affairs* (1988).

transmitter for 15 minutes. Doing so will interrupt the broadcast of the World Cup soccer match to millions of people. The match is scheduled to last for an hour more. If we wait that hour to free Jones, his injury will not get worse, but he will be in constant pain from both the mashed hand and continuous electrical shocks. “Does the right thing depend on how many people are watching—whether it is one million or five million or a hundred million,” asks Scanlon? Surely not, concludes Scanlon. We ought to free Jones now, even if the aggregate loss in utility to the millions of soccer fans far outweighs an hour’s additional pain to Jones: “if one can save a person from serious pain and injury at the cost of inconveniencing others or interfering with their amusement, then one must do so no matter how numerous those others may be.”⁴¹

But we trade off lives and greater and lesser harms all the time based solely on the numbers. Consider the following:

(iii) We have a limited supply of flu vaccine to distribute. If we give priority to the elderly and children, who are more likely to die from the flu, we will save many more lives. May we do so?

(iv) It is proposed that we build a new baseball stadium in the San Francisco Bay area to replace the dilapidated and freezing Candlestick Park. In past construction projects of this scope, on average two passersby have been seriously injured or died from falling debris. May we build the stadium?

⁴¹ What We Owe To Each Other, p. 235. For a similar conclusion, see Thomas Hill, Jr., “Basic Respect and Cultural Diversity” in *Respect, Pluralism and Justice* (2000), 78-79. Cf. Cranor, in *Lewens*, ed.

(v) Same as (iv), except that at a cost of \$50 billion dollars (approximately \$5,000 per taxpayer) we can add additional safety features that will reduce the expected injuries or deaths to one. Is it permissible to build the stadium without spending the additional \$50 billion?

If (as I would guess) most nonconsequentialists would answer yes to hypotheticals (iii) - (v), why do some balk at letting the numbers alone drive the results in cases like Taurek's (i) and virtually all balk in cases like Scanlon's (ii)? I believe the answer has nothing to do with the impermissibility of aggregation per se or of aggregating incommensurate harms. It lies instead in the presence of some or all of the confounding features above, which stir the sympathetic imagination of the authors and readers on behalf of would-be victims (identified victims, certain rather than statistical lives lost, harm imposed as a means to save others and/or through the deliberate and active agency of the chooser, the chooser's having to witness 'up close and personal' the consequences of the choice-- in Scanlon's case, to stand idly by while someone suffers before your eyes.)

Scanlon explicitly acknowledges this, when he compares the case for rescuing Jones with the case for never putting him in jeopardy to begin with (a variant on my hypo (iii)). Suppose we are considering building a new system of transmitting towers that will improve the quality of reception for many television viewers, and know that in the course of a project like this, some number of workers will likely suffer harms as great as Jones's. Scanlon concludes it is permissible to proceed with the project, notwithstanding that it "involve[s] risk of serious harm to others," provided the builders use adequate precautions to reduce that risk. As the

consequences are identical in the two scenarios, the consequences (alone) cannot explain why we are permitted to sacrifice the construction workers even as we save Jones. The explanation, says Scanlon, is that the construction case involves “failing to prevent accidental injuries rather than either intentionally inflicting serious harm on a few people or [as in Jones’s case] withholding aid from people who need it.”⁴²

Scanlon’s explanation, while sufficient to distinguish Jones’s case from the case of the construction worker, is surely incomplete. After all, we fail to provide aid daily to billions of people who need it more desperately than Jones, in order to satisfy desires at least as frivolous as the soccer fans’. As Elizabeth Ashford has argued, it is not clear why the steady state of chronic malnutrition that kills millions each year in underdeveloped countries doesn’t qualify as a constant, ongoing “emergency” subject to Scanlon’s Rescue Principle.⁴³ To explain (if not justify) why one would give precedence to Jones but not those billions, one would have to point to two additional features of Jones’s case: an identifiable victim, and harm that occurs ‘up close and personal,’ requiring us to stand idly by while Jones suffers should we choose not to rescue him.

⁴² Id. at 236. Consent also undoubtedly plays a role in this and other hypotheticals. Scanlon explicitly requires that, for the construction project to be permissible, the construction workers must consent to the risk they are assuming. Consent is surely morally relevant, but it differentiates Jones’s case from the construction workers only if we add in the requirement that the harm that ensues had to have been specifically foreseeable. And it doesn’t differentiate Jones’s case from Ashford’s rescue case at all. For an argument that neither consent or the difference between “failing to prevent harm” and “failing to aid” someone experiencing harm can bear the moral weight Scanlon gives it, see Norcross, “Contractualism and Aggregation.”

⁴³ Elizabeth Ashford, “Utilitarianism, Integrity and Partiality,” *J Phil* 97 (2000): 421-39. As Ashford notes, a narrow construction of emergencies covered by the Rescue Principle has obvious prudential advantages-- in particular, it avoids an interpretation of our duties to others that would be scarcely less demanding over a lifetime than what utilitarianism is thought to impose. Ashford, “The Demandingness of Scanlon’s Contractualism,” at 280, 287-88. The question here, however, is what descriptive differences in the cases account for the intuition that we can distinguish them morally.

III. To What Do The Consequences Matter?

But clearly when nonconsequentialists say that it is the consequences that matter, sometimes they really mean it. The question is, matter to *what*?

1. *The consequences determine whether we should have engaged in the conduct that produced them.*

As Thomas Nagel famously put it: “It is tempting ... to feel that some decision must be possible, in the light of what is known at the time, which will make reproach unsuitable no matter how things turn out. But this is not true; when someone [takes a justified gamble] he takes his life, *or his moral position*, into his hands, because how things turn out determines what he has done.”⁴⁴

But what exactly does it mean to say that “how things turn out determines what [one] has done”? The claim has been cashed out in a number of ways. I want to start with a central meaning it has been given in the nonconsequentialist literature, and the one of particular interest here: The consequences of an act determine in whole or in part whether the act itself was impermissible, in the sense that we would enjoin it from ever occurring if we could, costlessly.

Suppose we take at face value the claim that the consequences matter-- indeed, may often be dispositive-- in this sense. Suspending final judgment of the permissibility of an act until after its consequences are known has generated two sorts of paradoxes in the nonconsequentialist literature.

⁴⁴ Moral Luck, 29-30 (itals added).

The first results from treating bad consequences as *necessary* for conduct to be wrongful.⁴⁵ Within the legal academy, Ernest Weinrib has been the strongest defender of the view that actual injury is necessary for moral culpability.⁴⁶ But he is hardly alone in that view. Consider this from John Goldberg and Benjamin Zipursky: “The wronging of one person by another is the very essence of the enterprise, and until such an event happens, there is no occasion to inquire whether an actor can or should be held to have acted wrongfully by violating a moral or legal obligation of conduct.”⁴⁷ And this: “A claim for heightened risk—even if that risk is intelligible as a harm—does not invoke the sort of harm that defendants have a duty to take care to avoid causing.”⁴⁸ Or this, from Judith Thomson: “Suppose I play Russian roulette on you. (Gun with six chambers, one bullet.) And suppose that nothing happens: the bullet was not under the firing pin when I fired. Suppose I did this without your knowledge, so that you were caused no fear. Did I infringe a right of yours? It does not seem obvious that I did.”⁴⁹

⁴⁵ While there is not room to elaborate the argument here, I think this argument confuses the question of whether conduct is permissible with the question of whether any individual has reason or standing to complain about it. As noted above, in the typical tort suit, would-be plaintiffs don’t come forward until after they have been harmed, for the simple reason that they have no reason to know they are in danger until then. But where they do know they are in particular jeopardy, they may seek an injunction *ex ante*. More importantly, the fact that individual plaintiffs might not have standing to complain *ex ante* doesn’t preclude the state, acting on behalf of would-be victims, from prohibiting the risky conduct in the first instance. Every time we mandate safety standards by statute, regulation or common law, that is exactly what we are doing.

⁴⁶ See, e.g., The Idea of Private Law 153; Disintegration of Duty, 143.

⁴⁷ Goldberg and Zipursky, “Tort Law and Moral Luck,” at 1138. See also Goldberg and Zipursky, “Unrealized Torts,” at 1689); (“The duties typically recognized within the law of negligence are duties to take care not to cause “ultimate” ... harms, such as bodily harm or illness. By contrast, negligence law does not treat the ‘intermediate’ or ‘unripened’ harm of heightened risk as actionable injury.”)

⁴⁸ Goldberg and Zipursky, *id.* at 1634

⁴⁹ Get cite. See also Heidi Hurd, “The Deontology of Negligence,” *BU Law Review* 76: 249-72.

Translated to an action-guiding maxim, this view gives us no moral basis to intervene in conduct until it is too late to stop it, relegating the victim to what is (in most cases) a much inferior alternative: to seek compensation *ex post* for the completed harm. Not surprisingly, the view that rights theory leaves us helpless to stop rights infringements from occurring in the first place has been sufficiently unsettling to most nonconsequentialists to prompt a search for supplementary principles to explain the wrong of creating a risk of harm.⁵⁰

The second paradox, and the one of concern here, results from treating bad consequences as *sufficient* to establish the *prima facie*—or even conclusive—impermissibility of the act that produced them. Underlying this view is the intuition that somehow, if we had acted properly, our actions would not have resulted in harm to others. Herman’s Kantian account embeds that intuition in the maxim that the “end” we adopt and are responsible for bringing about must include no untoward consequences to others. Other contemporary rights theorists typically locate it in our duty “not to injure” or otherwise “produce an unwanted outcome.”⁵¹

⁵⁰ Stephen Perry, “Harm, History and Counterfactuals,” at 1307; Jules Coleman, *Risks and Wrongs*, at 198 (recheck cite); McKerlie, “Rights and Risk,” at 241; McCarthy, “Liability and Risk,” at 251; Dworkin, “Rights and Terror”; Samuel C. Wheeler III, “Self-Defense: Rights and Coerced Risk-Acceptance,” *Public Affairs Quarterly*, vol. 11, no. 4 (Oct. 1997) (arguing for the right to carry (and use) handguns for self-defense purposes, as consistent with the classic libertarian account of rights); Nozick, *Anarchy, State and Utopia*, ; Jean Hampton, “Correcting Harms versus Righting Wrongs: The Goal of Retribution,” 39 *UCLA Law Review*, 1661-1666 (1992); Heidi M. Hurd, “The Deontology of Negligence,” 76 *B.U. Law Rev.* 249, 262 (1996); Rahul Kumar, “Who Can Be Wronged?,” suggesting the wrong of risk is the “failure to have appropriately recognized [another’s] status as a person in one’s understanding of how it is appropriate to relate to her.” *P&PA* at 109; Claire Finkelstein (imposition of risk is itself a psychological injury); Christopher Schroeder, “Corrective Justice and Liability for Increasing Risks,” 37 *UCLA Law Rev.* 439 (1990); Kenneth W. Simons, “Corrective Justice and Liability for Risk-Creation: A Comment,” 38 *UCLA L.Rev.* 113 (1990); Christopher H. Schroeder, “Corrective Justice, Liability for Risks, and Tort Law,” 38 *UCLA L. Rev.* 143 (1990).

⁵¹ See, e.g., Judith Thomson, “Imposing Risks,” 173-74; Goldberg and Zipursky, “Unrealized Torts”; McCarthy, “Liability and Risk, at ; Weinrib, “Causation and Wrongdoing,” 63 *Chi-Kent L Rev* 407, 448-9 (1987); Ripstein and Zipursky, “Corrective Justice in the Age of Mass Torts,” in Postema, 218.

And contractualists like Scanlon and Kumar embed it in the supposition that a hypothetical person who discovers, peeking ahead, that she will be the one seriously hurt, killed or otherwise subject to consequences that “disvalue her humanity,” may for that reason alone have the right to “reasonably reject” the proposed conduct.⁵²

The view that we should act so as to ensure that we bring about favorable consequences is puzzling, given-- as Herman herself acknowledges-- that “[o]ur powers as agents are limited” and we know “from the outset that [we] cannot guarantee the success of [our] efforts.” How then are we to interpret “the idea that, even though our powers as agents are limited, our obligations are not just to ‘try’ to do what we ought to do,” but to succeed?⁵³

If we are to take it at face value, we are left with the central paradox posed by the problem of moral luck: We agree *ex ante* that it is permissible for you to drive down the street if you do so with due care (observing the speed limit, not tailgating, keeping your eye on the road, etc.), because on balance the activity is socially useful even if it poses an irreducible risk of harm to others, or because freedom of movement is necessary for people to pursue their own projects, or constitutes a right; but conclude *ex post*, after someone is hurt, that you acted wrongfully after all for reasons not fully captured by our *ex ante* decision rule.

That paradox is economically captured throughout the nonconsequentialist literature, when writers include the consequences of an act in the “act-description.” Thus, Judith Thomson describes moral theory as answering the question “what

⁵² Kumar, “Who Can be Wronged?”; Scanlon, *What We Owe to Each Other*.

⁵³ Herman, 98, 102.

harms we may or may not *cause* in what circumstance.”⁵⁴ John Goldberg and Benjamin Zipursky describe tort law as prohibiting “[t]he *doing of realized wrongs*,” not “the doing of unrealized. . . wrongs,” and as imposing a “*a duty to not injure*, rather than a duty to not engage in injurious conduct.”⁵⁵ Tony Honore includes in the category of “conduct” for which one may incur tort liability, “actions, omissions and *causing untoward consequences*.”⁵⁶

Where the consequences of acts are deemed to be certain *ex ante*, such formulations are probably best understood as a shorthand for “conduct that we can foresee *ex ante* will certainly cause harm X.” Thus interpreted, the command “do not harm others” serves as a non-paradoxical action-guiding norm (e.g., “Do not do act A, because it is certain to produce harm X”). But most (arguably all) acts entail consequences that are uncertain *ex ante*, and virtually all include at least a remote possibility of harm to others. As applied to that vast universe of conduct, there are

⁵⁴ Imposing Risks,” at 185 (italics added).

⁵⁵ “Unrealized Torts,” at . For other formulations of the same point, see *id.* at (“[N]egligence law in its main applications does not enjoin us to take care against risking harm to others. Rather, it enjoins us to take care against causing harm to others.”; *id.* at 1698 (it is a “very difficult question . . . whether the duty to be vigilant of causing a threat of HIV infection involves a duty not to cause *actual* exposure to HIV or a duty not to cause *possible* exposure through a medically possible means of transmission.”); Ripstein and Zipursky, “Corrective Justice in the Age of Mass Torts,” in Postema, 218, giving as an example of an unqualified duty not to injure “a duty not to cause another to die”; Julie Tannenbaum, “Emotional expressions of moral value,” *Philosophical Studies* (2007) 132:43-57, 47 (“one adopts the sub-end of not killing others...since this negatively impacts their rational capacity”).

⁵⁶ “The Morality of Tort Law—Questions and Answers,” in Owen, ed., *Philosophical Foundations of Tort Law* (Clarendon Press, 1995), 73-95 at 75. See also Owen at 98; McCarthy at 60: “Other things being equal, X ought not to cause Y’s death” ; Julia Driver, who conjoins “the actual deed” and “what is accomplished” in describing an externalist account of Kant (in distinction to “the internal state” of the actor). “Response to my Critics,” at 35; Lisa Heinzerling, “Knowing Killing” (arguing that actions that result in the death of others are always culpable, provided the risk was foreseeable). McCarthy gives the example of “murder” to show that an act-description can include both the *ex ante* intentions and behavior of an actor and the actual consequences of the act. (p. 57). But “murder” doesn’t describe an act; it describes the legal characterization of that act, which depends among other things on its factual consequences.

two ways to construe the injunction “do not injure others” to turn it into an action-guiding norm: “Do not do anything that might result in harm to others” or “Do not do anything that is certain to result in harm to others.” But the first interpretation rules out all action; the second rules out none (once we have, by hypothesis, removed harms that are certain to result). In either case, the prohibition is useless in helping us distinguish permissible from impermissible conduct.⁵⁷

The obvious difficulties with the argument that consequences determine the permissibility of the acts that produced them suggest that nonconsequentialists must mean something else when they say “the consequences matter.” But what?

But more often, when many nonconsequentialists say the consequences determine the permissibility of the conduct that produced them, they don’t really mean that the lorry driver should not have driven his bus, on his usual route and in his usual careful manner, that morning because of how things turned out. They mean something else. I turn to those possibilities now.

2. *The consequences matter because we experience them as mattering.* The most obvious way in which consequences matter is that we experience them as mattering, powerfully. A negligent pedestrian darts in front of your car, a suicide leaps out in front of a train you are driving, and dies. You know there is nothing, realistically, you could have done to prevent the death, but you nonetheless feel

⁵⁷ A third possible interpretation—“Do not carelessly injure others”—can serve as an action-guiding norm. But as noted below, that solution smuggles in (aggregative) consequentialist criteria through the back door.

terrible about your role in it, and others who also know on some level that it wasn't your fault nonetheless blame you.

It is hardly surprising that consequences matter in this way. We live in winners' (losers') history-- in what actually comes to pass, not in counterfactual states of the world that, *ex ante*, might have been much more likely to result from our actions. Only deeply disturbed individuals fail to register the real world they live in more strongly than the counterfactual worlds they don't, and fail to feel something akin to regret for harm they cause ("how much better if I had acted otherwise," in Williams's words).⁵⁸ As Strawson famously argued, our "reactive attitudes" to what actually happened and our causal role in it "have common roots in our human nature and our membership [in] human communities."⁵⁹

Inevitably, those reactive attitudes will color our retrospective judgments of others' and our own actions. Our reactive attitudes to bad outcomes are, as Strawson noted, ones we customarily associate with the moral domain: guilt and remorse for our own causal role, and outrage and reproach for others', in bringing about such harm. Not surprisingly, then, we *experience* those judgments as moral ones (he should never have left his wife, he should never have mailed that clock or carried it unprotected to the car, he should have driven the car more carefully). But that does not mean they *are* moral, if by moral we mean some more general principles about how we ought to live that can explain and justify our intuitions and

⁵⁸ "Moral Luck" at 27.

⁵⁹ "Freedom and Resentment," ; For other eloquent statements to the same effect, see Williams, Moral Luck, at 29; Railton, "Alienation, Consequentialism, Morality," in Scheffler, 99-100.

not merely reproduce them. It just means that morality does not have a monopoly on guilt and blame.⁶⁰

If one accepts that explanation, then there is nothing paradoxical about the situation Williams' lorry driver finds himself in.⁶¹ He made a choice *ex ante* to drive his usual route and in his usual (careful) manner, a decision that was prudent given all that was knowable at the time. *Ex post*, things turned out badly. The driver naturally regrets or feels guilty about what happened, and many will naturally blame him for it. His *ex ante* choice sounds in morality. If we conclude that from an *ex ante* perspective, given what he knew or reasonably could have known then, he made a morally justified choice, his *ex post* guilt with respect to that *ex ante* choice sounds in some other evaluative system entirely.⁶² That it is part of the human lot to choose wisely *ex ante* and yet feel regret or guilt and face blame *ex post* when things turn out badly is tragic.⁶³ But it is not a moral paradox.

Alternatively, many have taken the view that the lorry driver's *ex ante* choice and his *ex post* regret *both* sound in morality, but they implicate different moral

⁶⁰ Herbert Morris, "Nonmoral Guilt," at 221-222, in *Responsibility, Character, and the Emotions*" (Schoeman, ed.) (Cambridge 1987); Thomas Nagel, "Moral Luck," in *Mortal Questions*; Otsuka, "Incompatibilism and the Avoidability of Blame," *Ethics* (July 1998), 108:4, pp. 685-701, at 698-99. For an analysis locating our reactive attitudes to consequences in the psychological rather than moral domain, see Kahneman and Sunstein, "Indignation: Psychology, Politics and Law."

⁶¹ Williams, *Moral Luck*.

⁶² This claim pertains only to the harm-causing act; it does not speak to the question of whether there is a moral case for compensation *ex post*. On the latter, see point 3 below.

⁶³ See Martha Nussbaum, "Moral Limits of Cost Benefit Analysis, in Adler and Posner, . This paper focuses on one form of the tragic: bad consequences that result from reasonable choices *ex ante*, based on what was knowable *ex ante*. But there are many others that pose cognate problems, including tragedy in the classic sense, in which bad conduct is fated (Strawson's problem), or conduct that was freely chosen and that was or should have been known to be wrong *ex ante*, but turns out to have consequences that are so much worse than one could have reasonably expected to follow from relatively trivial badness. On the last, see Waldron, "Moments of Carelessness and Massive Loss."

issues.⁶⁴ Put in Kantian terms, the driver's ex ante choice implicates his obligation to adopt and pursue moral ends and sub-ends (means rationally calculated to achieve those ends). This the driver has done, when he drives down the street with the belief that he will not kill anyone and in a fashion (i.e., prudently) that he reasonably believes will ensure that result. The second implicates the moral value we attach to achieving our ends and sub-ends. This the driver fails to do when he causes a child's death. The regret he feels for that failure is regret for failing to bring about a state of affairs that "embod[ies] his moral values," which include "not negatively impacting" any individual's "rational capacity."⁶⁵

Whether it makes sense to describe as "moral" the regret we feel for the consequences of actions *no one would have had us perform differently* I leave to others. But it is not paradoxical, as such regret is understood to imply nothing about the permissibility of the action itself-- indeed, in contrast to the *guilt* we feel when the child dies as a result of our *negligence*, it presupposes the moral rightness of the act.⁶⁶

Finally, some have taken our reactive attitudes as evidence that the lorry driver's ex ante choice was impermissible after all. On its face, this position lands us back in consequences mattering in sense (1)-- back, that is, in the paradox of moral luck. But often, the only operational implication of an author's embracing that paradox is to reassure all of us we are morally justified in reproaching the lorry driver ex post for acts we all would have agreed were morally permissible ex ante.

⁶⁴ For an exceptionally clear explication of this viewpoint, see Tannenbaum, "Emotional expressions of moral value."

⁶⁵ Tannenbaum, 47, 55.

⁶⁶ See Tannenbaum at .

G.E. Moore's distinction between justifiable and permissible actions (or subjective and objective permissibility) seems to me to fall into this category. Moore argues that whether actions are justifiable (subjectively permissible) depends at least in part on what the actor (reasonably???) believes will be the consequences. Whether they are objectively permissible depends on the actual consequences.⁶⁷ But the only thing that appears to follow from concluding that something is objectively impermissible is that we are justified in so describing it. If so, we are back to the reality that talk is cheap. When talk is all that is at stake, we can live with the paradox created by our temporally inconsistent judgments, and do so all the time.⁶⁸

3. We are morally required to respond to the consequences of our actions.

Often, when nonconsequentialists say that "how things turn out determines what [someone] has done," what they really mean is that they determine whether the actor owes the victim compensation or an apology or some other response *ex post*.

The state has two decisions to make in regulating potentially harmful conduct: (a) what risks to allow citizens (or the state, acting collectively) to impose

⁶⁷ Moore, *Ethics* (1958), at 118-121 (recheck cite).

⁶⁸ For what it's worth, I believe that the same "talk is cheap" principle explains why Strawson, in his famous article, can leave unresolved the apparent moral inconsistency between a strong version of determinism and our reactive attitudes of guilt and blame. Since all that is at stake in resolving it is how we are going to characterize the reactive attitudes we all agree we are going to continue to display whether justified or not, Strawson has the luxury to refuse to answer. But suppose Strawson were asked by the state to design a system of punishment that tracked true moral desert, excusing all those wrongdoers who truly "could not have done otherwise" than they did. Now, alas, some decisions must be made.

on others; and (b) who should bear the resulting costs.⁶⁹ (Typically such costs will be incurred in the form of harm resulting from the act. But they can also be incurred in the form of economic losses suffered by the would-be actor as a result of the state's prohibiting the act as too risky for others.) In the moral realm, we face essentially the same two questions, except that the range of available responses in (b) is broader, including (in addition to monetary compensation) apology or replacement of lost goods.

The decisions required by (a) and (b) are very different. In economic terms, the first poses an allocative question: what conduct ought we to permit and what prohibit in a just society? The second poses a distributive one: how do we distribute the costs that result from permitting or prohibiting a given act? Both decisions are important, but they are different, and implicate different considerations. For the consequentialist, the answer to (a) —should we permit X act-- turns on whether permitting it will produce more (social) good than (social) harm, all things considered. The answer to (b) —who should bear the costs of any resulting harm-- turns on a host of factors, including which result will produce the optimal incentives for all the parties to act in the socially optimal fashion going forward, and which party can best absorb the loss (that is, at the lowest loss in welfare).

For the nonconsequentialist, the answer to (a) presumably turns on whether it is impermissible to do X, all things considered (because doing so will violate

⁶⁹ The distinction between the permissibility of conduct and the distribution of costs that arise from prohibiting or allowing it has been familiar to legal audiences at least since Calabresi and Melamed's famous article spelling out the combinations that result from possible answers to the two questions. "Property, Liability and Inalienability: Three Views of the Cathedral," 85 Harv. L. Rev. 1089 (1972).

someone's rights, because it treats the victim as a means rather than an end, because the victim could have reasonably refused to have consented to its being done). The answer to (b)— how to distribute the losses associated with X's occurrence or prohibition-- turns on a different question: whom is it fair to ask to bear those losses? These needn't be answered the same way. We can conclude, all things considered, that it is better to allow Smith to do X than to prohibit it, even though we foresee that doing it will result in harm to some individuals, but that it is fair to require either the state (through a social insurance scheme) or Smith (via a strict liability rule) to compensate those individuals who are harmed.⁷⁰

Conversely, we can decide to prohibit Smith from doing X (because, for example, serious harm to others is foreseeable if X occurs), but require that either those individuals benefited by the prohibition on X or (more likely) the state should compensate Smith for the loss entailed in not being able to do X.⁷¹ Of course, how the state resolves the distributive question will indirectly affect the allocative decision by affecting the incentives Smith has to engage in act X in the first place. But (a) and (b) are nonetheless separate questions, both analytically and operationally.

On occasion, nonconsequentialists have explicitly tried to address both questions, but in conflating the two, have answered neither. (Nozick's treatment of

⁷⁰ For a suggestion that this may frequently be the appropriate response to tragic choices, see Martha Nussbaum, "The costs of tragedy: some moral limits of cost-benefit analysis," in Adler and Posner, eds., *Cost-benefit Analysis* (pp. 170, 173, 177-78, 180); Frances Kamm, *Intricate Ethics*, p. 240.

⁷¹ Takings law would be an example of the latter. For an example from common-law torts, see *Spur v. Del Webb*.

the problem in his chapter on “Prohibition, Compensation and Risk” offers a particularly spectacular train wreck along these lines.) More often, I believe, what is really bothering nonconsequentialists is (b)--- the obligation to respond in some fashion to the harm one causes others. Some authors have been explicit about this.⁷² In other cases, it is implicit in the account of what is morally required of the actor. Consider this statement from Barbara Herman: “[A]lthough good intentions plus adequate care are enough (whatever the outcome, the agent has done what she ought), they are not the end of the story (*the agent may need to respond to what actually happens when she acts*). There is room for consequences, actual consequences, to make a moral difference....”⁷³

I believe it is much harder than nonconsequentialists think to explain why *fairness* (as opposed to efficiency considerations) requires an actor to compensate the victim for the consequences of an act we told him it was perfectly permissible to engage in.⁷⁴ But requiring him to do so is not paradoxical, in the sense of providing contradictory answers to the identical question (was the act in fact permissible?).

⁷² I have in mind here Ernest Weinrib, Jules Coleman, Stephen Perry, Martin Stone and others who have focused their efforts on showing that once we have established harm-causing behavior to be tortious by other (unstated) standards, the US tort system, by requiring the tortfeasor rather than (say) the state to compensate for harms, adopts a corrective justice rather than welfarist approach. See, e.g., Stone, “Doing and Allowing”; Coleman, “Tort Law and Tort Theory,” in Postema, 183.

⁷³ Id. at 98 (itals. added). See also id. at 98: “The basic idea is that the circumstances of failure call for a moral response.”

⁷⁴ The difficulties with this position have been treated at length in the legal literature. In brief: (i) A true strict liability regime is a conceptual impossibility for Coasian reasons. In all “strict liability” cases, some prior decision has been made on normative grounds to hold one actor or the other causally responsible for the joint costs of their conflicting activities. (ii) More importantly, assuming nonconsequentialists can articulate morally compelling criteria that do not themselves boil down to some form of cost/benefit calculus for assigning causal responsibility as between joint causes, if the party to whom we assign causal responsibility acted as we would have him act, why does fairness require that he compensate the victim? Why is his case any different, morally speaking, from the case where harm results from (say) an act of nature?

4. Consequences matter because of the resulting plight of the victim.

Pedestrian Abel darts out in front of Baker's car, and Baker, through no fault of her own, hits Abel. To most people, the plight of Abel-- who is now (let us suppose) seriously disabled and too poor to pay for medical treatment-- matters, in the sense that it imposes a moral imperative on others to help Abel out.⁷⁵ This obligation to help Abel just because he is suffering is not agent-relative. It responds to Abel's urgent need, not the circumstances that produced it. To put it another way, it sounds in distributive rather than corrective justice. Under this view of why consequences matter, in theory Baker has no greater obligation to aid Abel than would an innocent bystander or someone living 3000 miles away. Not surprisingly, however, the humanitarian reasons for everyone (including Baker) to aid Abel are hard to keep from bleeding into the intuition that Baker in particular owes Abel something in virtue of his role in causing the harm, *even if his conduct was permissible, indeed above reproach.*

Often, when authors purport to pin liability on causation simpliciter, they are implicitly assuming the actor to have been at fault (acted negligently, etc.). Shelly Kagan's classic defense of strict liability in "Causation and Responsibility," *Am Phil Quarterly* 25:4, 293-302 (1988), is a case in point. While Kagan purports to be addressing the claim that the mere fact of "having caused harm generates a special obligation to aid the victim of that harm" (id. at 293, italics added), he states he is going to "sidestep" the more complicated case of the "[a]gent [who] may *not be at fault* for having caused the harm" (id., italics added). I'm not sure what Kagan has in mind here, but at least on its face the statement seems to limit the scope of Kagan's argument to conduct that was negligent or otherwise clearly wrongful under conventional notions of fault. His central example—the obligation of Agnes, who pushed Victor into the pool, to rescue him and pay for any resulting expenses he incurs-- hardly allays the concern that the work Kagan thinks is being done by the mere "causing of harm" (id. at 300) is really being done by the fact that the harm-causing conduct was negligent or worse. See also Judith Jarvis Thomson, "Remarks on Causation and Liability," *Phil. And Pub. Affairs*, vol. 13 (1984), pp. 101-113., [recheck Thomson]

⁷⁵ See, e.g., Owen, "The Fault Pit," 26 *Ga L Rev* at 716: "By definition, an accident diminishes the quality (and perhaps quantity) of a victim's life and other goods, which produces suffering for the victim. We may assume that human suffering is undesirable and so should be avoided, ex ante, or remedied, ex post."

5. *Consequences may matter because they are probative of negligence.*

Finally, consequences may provide relevant evidence as to whether the defendant acted negligently.⁷⁶ A crane topples from the top of a tall building. Other causal explanations may ultimately surface (the crane was defective, there was an unprecedented wind), but the mere fact of the crane's falling strongly suggests the operator did not use due care in securing it to the top of the building. At the extreme, when we cannot imagine any explanation for the consequences other than the carelessness of the actor, the law treats outcomes as sufficiently probative of negligence to shift the burden of proof over to the defendant (the doctrine of *res ipsa loquitur*).

Of course, a significant danger exists that hindsight bias will lead us systematically to attribute more evidentiary weight to consequences than they merit. But in principle, there is nothing paradoxical about saying consequences may matter in this sense.

While all four of these alternative explanations for why the consequence matter avoid the paradox of moral luck, like explanation (1), they offer no clear alternative to aggregation for deciding which sorts of risks it is permissible to impose on others and which it is not.

⁷⁶ See Herman, *id.* at 107, noting that failure, at the extreme, may "cause us to suspect" gross negligence or even the sincerity of the original promise. Strawson, *Freedom and Resentment*, at : "The fact of injury creates a prima facie appearance of [the] demand [of inter-personal regard's] being flouted." Otsuka; Latus, "Moral and Epistemic Luck," *J of Philosophical Research* 25:149-72 (200); N. Richards, "Luck and Desert," *Mind*, 65 198-209 (1986); Thomson, "Morality and Bad Luck," in *Moral Luck* (Statman, ed), 1993. Williams may have had some such concern in mind in singling out "intrinsic" (as compared to "extrinsic") luck as indicative of moral failure. "Moral Luck," 25-27.

E. So what?

As I see it, the nonconsequentialist literature on permissible harm to others has left two questions largely unaddressed. The first is the scope of our duty to rescue those in immediate, serious need. The second, and the focus of this paper, is the relevance of our “duty not to harm others” to garden-variety decisionmaking under uncertainty.

These are not small omissions. For every Baby Jessica trapped in a well or technician trapped in a broadcast booth, there are hundreds of thousands of other people facing certain death every day whom we have the means to rescue. They are real, identifiable people; we just don’t happen to know who they are or care about their plight. Does our duty to rescue end with the Baby Jessicas? There are powerful psychological and emotional explanations for why we might draw the line at the Baby Jessicas. There are economic ones as well. (As long as the duty to rescue is limited to the Baby Jessicas of the world, we can maintain the illusion of nonscarcity that underlies our willingness to spend millions to save one life. But when we start scaling up the duty to rescue, we come up against the real scarcity of resources very quickly.) The question is, are there moral reasons to draw the line there?

Likewise, for every real-life variant on the trolley problem or Jim and the Indians, millions of decisions get made every day that pose uncertain risks of potentially serious harm to unidentified others. Every time we get in a car, market a new drug, decide whether or not to recall a product, we are taking others’ lives into our hands.

I'm not sure anything practical turns on resolving the first question—the scope of our duty to rescue. For whatever reason— psychological, emotional, moral—there is no possibility we are going to stand idly by and let Baby Jessica perish in the well. But there is also no political will at present to expand our collective “duty to rescue” beyond these high-visibility, up-close-and-personal, cases.

But a lot turns on the answer to the second. Perhaps not in the realm of individual decisionmaking, where the one-off nature of most of our choices and our limited capacity for ex ante rational decisionmaking both minimize the impact of the moral luck paradox on our ex ante choices and ex post regret.⁷⁷ But surely in the public policy realm. There, we are typically choosing rules that will govern thousands or millions of events over the long term. As a result, even very low-probability harms are overwhelmingly likely to occur at some point. If we parole enough prisoners, one of them will turn out to be Willie Horton. If we develop a vaccine for AIDS that has a one out of a million chance of triggering a fatal reaction and administer it to enough people, someone is going to die from it. If the world's most conscientious child welfare agency has to supervise thousands of children from dysfunctional families, over the long run some child under its watch is going to die from abuse or neglect.

⁷⁷ Among other things, because most individual choices are one-off decisions, our inability to go back in time limits the practical influence of guilt for past sins on our behavior. In addition, since in the event most choices do *not* result in serious harm to others, we have fewer occasions to be consumed by regret or blame ex post for bad outcomes.

How we understand and respond to those bad outcomes when they inevitably come to pass is of enormous social importance. We all know how the average citizen responds: You (government officials) had a duty not to harm anyone, or a duty to keep us safe, and you violated that duty. (Or, as one commentator dryly remarked, "Seldom do we hear a company that was responsible for a deadly accident justify the loss of lives by saying that it was the result of a decision which, in terms of its [expected??] effects, produced far more good than harm."⁷⁸) And politicians and policy makers know how the average citizen will respond as well. The result is to drive politicians and other public employees to channel enormous resources into preventing the high-visibility bad outcomes for which they know they will be held responsible. Negotiations underway right now to ease California's budget crisis and prison overcrowding by releasing nonviolent offenders early are hung up on exactly this problem. As one Republican legislator acknowledged with refreshing candor: "If we let someone out early, and that man commits a crime, the Assembly members are worried that that will come back to haunt them like the old famous Willie Horton ads."⁷⁹ Social workers report informally that child welfare agencies have responded to the recent high visibility deaths of children in foster care by directing almost all of their resources to children they think might be at some risk of death, ignoring the thousands of others that are facing serious but (they believe) nonlethal threats of abuse and neglect.

⁷⁸ Hansson, "Risk and Ethics," in Lewens, *Risk: Philosophical Perspectives*, 32.

⁷⁹ <http://www.nytimes.com/2009/08/31/us/31abduct.html>

The same risk aversion about ex post blame in the public and private sector drives a substantial portion of our health care expenditures, investments in “homeland security,” and countless other major public policy decisions. It may be that our propensity to think that bad outcomes imply bad actions is immune to reason, and the only way policymakers can curb its influence is to manipulate the psychological salience of bad outcomes. But surely, before we reach that question, we should answer the question whether it *ought* to be curbed-- whether, that is, the consequences of our actions do in fact determine, in whole or in part, the permissibility of those actions.

The full-blown consequentialist will say, of course they don't. It is tragic when reasonable actions have bad consequences, but it is nobody's fault. In a world of scarcity (in the broad sense), whatever we do has potential costs to someone. The best we can do is to act in the way that minimizes expected costs relative to expected benefits, and, if the costs to individual victims are serious enough, to remediate them ex post, on welfarist grounds.

Some nonconsequentialists have said, in effect, you are right. Nonconsequentialist principles can tell us what it is permissible to do when the conduct is *malum in se*, or when we know with certainty the harm that will result. But we concede they cannot handle the garden-variety problems of decisionmaking under uncertainty, and that such problems will have to be resolved in some other way—most likely some aggregative function.⁸⁰ Others have demurred, leaving the

⁸⁰ See, e.g., Heidi Hurd, “The Deontology of Negligence,” *BU Law Review* 76: 249-72; Nozick, *Anarchy, State and Utopia* at 28, 149 (proposing a “utilitarianism of rights”); McKerlie, at 240-4;

problem of decisionmaking under uncertainty to another day, or have qualified the “duty not to harm” in ways that appear to turn it into some version of a cost/benefit calculus.⁸¹ The challenge for both groups, it seems to me, is to show that, if nonconsequentialist principles are irrelevant to garden-variety decisionmaking under uncertainty, they have any significant role to play in regulating conduct (at least outside of the criminal context.)

Hardcore nonconsequentialists, if they are to be taken at their word, have stuck to their guns, insisting that the prima facie right not to be harmed extends to all conduct, including decisionmaking under uncertainty, imposing in all cases a correlative duty on actors “not to harm others.” That duty is not just to use reasonable efforts not to harm others; it is to succeed. It therefore rules out aggregation in any form, as we may not put some individuals at risk of substantial harm in order to realize expected greater benefits for others.

The challenge for this last group of nonconsequentialists is different and more serious: to articulate the alternative criteria (inspired by deontic,

Coleman; Goldberg and Zipursky; Judith Thomson, “Imposing Risks,” at 185; Owen, “The Fault Pit,” at 716.

⁸¹ Some proposed alternative criteria seem indistinguishable from a cost/benefit calculus. See, e.g., Joel Feinberg, Harm to Others, at 190-93 (culpability for harm turns on three factors: the probability of harm, the magnitude of expected harm, and the independent value of the risky act to the actor himself and society at large); Charles Fried, Right and Wrong, at 12 (it is impermissible to expose others to “an undue risk of harm, but what risk is undue is a function of the good to be attained and the likelihood and magnitude of harm”). Others are more ambiguous, but as currently articulated do not rule out the possibility that they will cash out into some form of cost/benefit calculus. Thomson; Wheeler, *id.* at 432-33 (whether a given right “survives” its encounter with a conflicting right “depends on the relative strength of the rights”); McCarthy, “Actions, Beliefs,” at 65; Rahul Kumar, “Who Can Be Wronged,” at 107; Tannenbaum, “Emotional expressions of moral value,” at 51 (one’s “obligation not to kill” does not extend to situations in which, “through no fault of one’s own, [one] could not avoid killing...in the way that [one] did); Dworkin, “What are Human Rights?,” pp. 39-40.

contractualist or other nonconsequentialist principles) by which we should determine what forms of potentially harmful conduct to permit and what to disallow, in our individual and (even more importantly) in our collective decision-making capacities.

If there is an answer to this, it needs to be put on the table so that it can become part of the public debate over risk regulation and assessed as a credible substantive alternative to some version of cost-benefit calculus or other aggregative criteria. If there is not, then nonconsequentialists have an obligation to acknowledge that, rather than lending the imprimatur of morality to what amounts to nothing more than hindsight bias. We may not be able to change the widespread intuition that bad consequences imply bad conduct; it may be an irreducible part of what it means to be human, and need to be accommodated in some fashion in our public policies simply in virtue of that fact. But if the intuition is wrong, philosophers ought to be the ones saying so most clearly, and doing what they can to counteract the pernicious public-policy consequences of all of us at least half-believing otherwise.