NEGLIGENCE*

By Kenneth W. Simons

I. Introduction

Faced with the choice between creating a risk of harm and taking a precaution against that risk, should I take the precaution? Does the proper analysis of this trade-off require a maximizing, utilitarian approach? If not, how does one properly analyze the trade-off?

These questions are important, for we often are uncertain about the effects of our actions. Accordingly, we often must consider whether our actions create an unreasonable risk of injury—that is, whether our actions are negligent.

Consider two examples:

(1) *The (mythical) Ford Pinto:*[1] The manufacturer of an automobile discovers that strengthening the fuel tank on 12.5 million existing vehicles would cost $11 per vehicle and would prevent 180 burn deaths, 180 serious burn injuries, and 2,100 burned vehicles. Calculating a unit cost of $200,000 per death, $67,000 per injury, and $700 per burned vehicle, the manufacturer concludes that the total cost of preventing the injuries is $137.5 million, while the accident losses that the precaution would prevent amount to $49.5 million. Accordingly, the manufacturer chooses not to take the precaution.

*I thank participants at the Boston University School of Law faculty workshop, as well as the other contributors to this volume, and its editors, for their helpful questions and comments. David Schur provided valuable research and editing assistance. I am especially indebted to Hugh Baxter, David Lyons, and Larry Solum for their advice.

[1] See Gary T. Schwartz, "The Myth of the Ford Pinto Case," *Rutgers Law Review* 43, no. 4 (Summer 1991): 1013–68, for a full account. See also Richard A. Posner, *Tort Law: Cases and Economic Analysis* (Boston: Little, Brown & Co., 1982), 225–26.

The account in the text reflects the popular mythical account of the Ford Pinto controversy, but it is misleading in critical respects. On the one hand, it ignores the question of Ford's responsibility and liability for an original negligent design in locating the fuel tank in an unusually vulnerable position. Of course, if the original design was defective, Ford should be liable for any resulting injuries, even if the post-manufacture precaution of strengthening the fuel tank was too costly. On the other hand, the dollar values were computed in 1973; Ford did not, in fact, rely on these figures in deciding against strengthening the fuel tank; the controversially low "value of life" figures were supplied by a federal government agency, not by Ford; the vehicles included all cars sold in the United States in a typical year, not just the Ford Pinto; and the study in question focused on rollover accidents, not rear-end collisions. See Schwartz, "Myth," 1020–28.

(2) *Two speeding drivers*: Amy drives at high speed to the hospital to obtain medical care for her child, whom she reasonably believes to be in need of emergency medical care. Beatrice drives at high speed to a critical business meeting; she reasonably believes that if she misses the meeting, there will be a significant delay in implementing a health delivery system, a delay that might cost several lives.

For many, the mythical Ford Pinto example epitomizes the defects of a utilitarian approach to negligence, especially the cost-benefit variant of that approach. And the comparison of Amy with Beatrice creates similar worries. Amy, but not Beatrice, seems justified in speeding; but this suggests that whether speeding is justified does not depend merely on the level of benefits that speeding would foreseeably produce.

This essay examines the question of responsibility for negligence mainly from the perspective of private morality, though it also analyzes legal norms embodying prohibitions against negligence. I hope to illuminate the complexity and richness of a problem that is often treated in reductionist fashion, not only by maximizing, utilitarian approaches, but also by some leading deontological approaches.[2] I also will suggest that negligence is better understood as an aspect of fault than as an aspect of corrective justice, contrary to the prevailing deontological views.

In Section II, I more carefully define the problem of negligence. First, I differentiate three senses of negligence (unjustifiable risk, conduct that violates a "reasonable person" criterion, and culpable inadvertence). This essay focuses on the first sense. Second, I emphasize that negligence presupposes an *ex ante* perspective. A negligent actor is one who either realizes, or should realize, that she has (unjustifiably) created a "low-probability" risk of harm (in a sense that will be explained). Third, I ask what is special about negligence. Do distinctive moral principles apply, or is negligence simply an instance of moral principles that would normally apply *ex post*, or *ex ante*, if we knew to a certainty the results of our actions? For the most part, I conclude, negligence is not a qualitatively distinct subject of moral inquiry.

In Section III, I explore negligence as an aspect of private morality, of personal responsibility: What should a person do when faced with a choice between risk and precaution? The common-sense moral precept that one should not be negligent, I conclude, reflects neither the coldly calculating utilitarian conception suggested by some forms of economic analysis of law, nor an absolutist deontological conception that blithely ignores the consequences, costs, or disadvantages of taking precautions

[2] Broadly speaking, utilitarian approaches judge the morality of an act by the aggregate utility of the consequences of that act, while deontological approaches instead (or also) consider whether the act is right or wrong in itself.

against risk. Beginning with a utilitarian account, I progressively modify the analysis to encompass a variety of nonutilitarian concerns. In the end, a pluralistic balancing approach is the most suitable for recognizing the breadth of values expressed in ordinary moral judgments about risk and in relevant nonutilitarian principles.

In Section IV, I briefly examine some distinctive features of law and analyze more carefully how legal norms of negligence should be defined and enforced, with particular attention to their relation to private moral norms.

Finally, in Section V, I suggest that principles of fault (or unjustified conduct), rather than of corrective justice (or the correction of harms), are the better interpretation and more convincing deontological justification of Anglo-American tort doctrine. The *ex ante* fault perspective supports a primary duty not to *act* negligently; the duty to compensate for harms one has negligently caused is distinctly secondary. Although compensation is ordinarily the only feasible remedy for isolated acts of negligence, this is a contingent fact, not a necessary implication of the deontological view of negligence.

II. The Scope of the Problem

A. What is negligence?

The topic of this essay is "negligence," by which I mean the failure to take a reasonable precaution against risks of harm. A "negligent" actor is one who acts as he should not have acted (or omits to act when he should have acted), and thereby creates an unreasonable risk of injury to others. This meaning of negligence has great importance in both morality and law, as I will try to show.

Negligence has other meanings. Negligence can characterize *beliefs*, rather than conduct. Thus, a person's beliefs might be negligent, meaning that her subjective conviction (that X is the case) is not based on reasonable grounds. Or she might be negligent in *not* believing Y (meaning that she should believe Y, or that a reasonable person would believe Y).[3] Moreover, a person might express a negligent *attitude* toward her conduct or the results of her conduct—meaning, perhaps, that she failed to show a reasonable degree of concern. Further, negligence can refer only to some aspect or aspects of a person's conduct, beliefs, or attitude, not to a global judgment about how she should have acted. We might conclude that a driver was negligent in not noticing a pedestrian, without necessarily

[3] However, the question of whether an actor's beliefs are reasonable is, in a limited way, relevant to whether the actor's conduct is reasonable. For we cannot make sense of the concept of an *ex ante* probability of a risk of harm without an epistemic account of risk. See Section IIB below.

implying that he must have been negligent and should have taken a precaution against causing an accident, *all things considered*.[4]

Three conceptions of negligence are especially important in morality and law. One conception emphasizes *unjustifiable risk*. A negligent act is one that creates an unreasonable or unjustifiable risk of future harm. A second emphasizes evaluation according to a *"reasonable person" criterion*. A negligent act, belief, or attitude is an act that a reasonable person would not perform, or a belief or attitude that a reasonable person would not harbor. A third conception emphasizes culpable or unreasonable *inadvertence*: the actor, although not consciously aware of a risk, should have been aware.[5] This conception is often employed to distinguish the negligent actor from the "reckless" actor who recognizes an unreasonable risk before taking it.

The three conceptions often overlap in fact, but they are distinct in principle. One might call an act negligent because it creates unjustifiable risks, apart from whether a "reasonable person" would act differently. (Under a so-called "subjective" test of negligence, you are not negligent if you do the best that you can given your personal capacities, but you still might create unjustifiable risks.) Conversely, one might employ a reasonable person test for evaluating choices that have virtually certain consequences. (In self-defense, for example, the predominant legal test essentially asks whether a reasonable person in the shoes of the defendant would

[4] Suppose no precaution would have avoided the accident, even if the driver *had* been paying proper attention.

In law, negligence is sometimes a culpability requirement of only one element of a crime, tort, or other legal norm, in which case it might have subsidiary importance. In the crime of assaulting a police officer, for example, liability might depend on the actor being at least "negligent" as to the risk that the person he is assaulting is a police officer; but the more serious culpability obviously is that entailed by the act and intention of assaulting a person.

[5] As a matter of ordinary language, "negligence" might indeed presuppose inadvertence:

> "Carelessness" and, consequently, "negligence" signify neither a "state of mind," such as indifference, nor merely a "type of conduct," such as a failure to take precautions against harm. "Carelessness" or "negligence" is a failure to give active measure-taking attention to the risks inherent in the successful prosecution of some activity.

Alan R. White, *Grounds of Liability* (Oxford: Clarendon Press, 1985), 102.

However, I use the term in the broader sense that White calls a "type of conduct." The ordinary language usage is, I believe, beginning to expand toward this broader usage. In any event, the sense of negligence as unreasonably risky conduct has more general importance in morality and law.

Of course, inadvertence is not always culpable. Moreover, when inadvertence is culpable, this might be because one should have adopted a different action-guiding strategy that would have avoided risk, not simply because one "should have been aware" of the risk. See Joel Feinberg, "Sua Culpa," in Feinberg, *Doing and Deserving: Essays in the Theory of Responsibility* (Princeton: Princeton University Press, 1970), 194 ("Overly attentive drivers with the strongest scruples and the best intentions can drive as negligently as inattentive drivers and, indeed, a good deal more negligently than experienced drivers of strong and reliable habits who rely on those habits while daydreaming. . . ."); see also Kenneth W. Simons, "Rethinking Mental States," *Boston University Law Review* 72, no. 3 (May 1992): 550–51.

employ the intentional force that the defendant employed.)[6] Similarly, one might call a belief negligent because a reasonable person would believe otherwise; but this explanation need not presuppose that the actor has unjustifiably *risked* anything.

The first conception, the unreasonable creation of (and failure to take a precaution against) a future risk of harm, is probably the most important sense of negligence. The second, "reasonable person" conception is essentially a legal conception, reflecting certain pragmatic and institutional features distinctive of law, as I suggest below. The third, "inadvertence" conception is not without importance, but the problems that it presents are beyond this essay's scope.

An additional point about scope concerns the interests that a norm against negligent conduct protects. In both law and morality, negligent creation of a risk of *physical injury* has special importance. Accordingly, I restrict my analysis here to the interest in avoiding physical injury.[7]

A final point about the negligence concept is this: Lack of justification is part of the ordinary meaning of negligence. This feature creates an interesting asymmetry between negligence and the other forms of wrongdoing with which negligence is usually compared. Negligence is a composite concept: a negligent actor both creates a risk (that he could have avoided) and is unjustified in doing so. By contrast, an intentional killer intends to bring about a death, and a knowing killer (as conventionally defined) is one who believes that death is a virtually certain result of his acts; but in either case, it is possible that the actor *is* justified (for example, because he is defending himself against a culpable aggressor). Lack of justification is built into the very concept of negligence, but it is not built into the concept of intentionally or knowingly bringing about a harm.[8]

This asymmetry could be eliminated either by building lack of justification into the definition of intentional and knowing harms, or by isolating lack of justification as a separate element of negligence. On the latter approach, negligence could be divided into two parts: the creation of

[6] This is an oversimplification. The law typically predefines certain categories of force as conclusively reasonable, and adds an explicit "reasonable person" criterion only for certain questions, including the actor's belief that force was being threatened or that defensive force was immediately necessary.

[7] To be sure, sometimes "negligence" is used in a wider sense. One can negligently forget a spouse's birthday, or make a negligent accounting mistake that causes only economic harm. And the epistemic sense of "negligence" is very wide: with respect to any subject whatsoever, a belief can be formed negligently (i.e., without sufficient grounds), or one can be negligent in failing to form a true belief (based on the grounds available). In law, however, the most important sense of negligence is with respect to risks of physical injury, for those are the dominant uses of negligence in tort and criminal law, which in turn are the dominant legal fields in which the negligence concept is used. In morality, as well, negligent creation of physical harms has special importance.

[8] However, certain descriptions of intentional harms do presuppose that the harms are unjustifiable. To "murder" another is not merely to cause his death intentionally, but also to do so without justification. I thank Larry Solum for this point.

(say) a "significant" risk to others, and the lack of justification for creating "significant" risks to others. In the end, however, we would need to create a range of norms governing risk-creation—a norm governing "trivial" risk-creation (under which we inquire whether the actor had a justification for creating a "trivial" risk), a similar norm governing "modest" risks that do not reach the level of "significant," and so forth. But it is simpler to state the requirement this way: The justification for imposing a risk must, in general, be stronger as the probability and seriousness of the risk increase.[9]

Still, building lack of justification into the very concept of negligence has this important implication: a judgment that an actor is negligent is an "all things considered" judgment that the actor was unjustified and should not have acted as he did. By contrast, a judgment that the actor intentionally or knowingly caused a harm has no such implication.

In sum, this essay considers what one should do, when faced with a choice between risk and precaution. It examines negligence as an aspect of action-guiding morality, and as a legal norm expressing that moral norm. Of course, distinct and important questions remain with respect to the kind of person one should be, how one should express appropriate concern about the negative consequences one justifiably brings about (including feeling regret, apologizing, or compensating the victim), and similar difficult topics.

B. The ex ante *perspective*

Implicit in negligence analysis is a crucial assumption: that personal responsibility is judged *ex ante*, not (merely) *ex post*. Suppose we are trying to decide whether someone has acted negligently—for example, by driving 40 miles per hour around a particular curve under particular road conditions, or by performing an operation using one medical technique rather than another. Then we should imagine ourselves "stopping the videotape" (so to speak) at the moment when the decision to act was made, and determining then and there whether the actor should have acted differently, in light of the comparative risks and other reasonably expected advantages and disadvantages of the alternative action. A negligent agent is one who acts as she should not have acted, judged from this *ex ante* perspective.

[9] In criminal law, when the probability and seriousness of the risk are sufficiently great (e.g., when one knowingly or intentionally kills), the burden of persuading the fact-finder of lack of justification sometimes shifts to the defendant, and the grounds of justification are also limited to certain narrow categories such as self-defense, defense of others, or necessity. These legal features reflect the fact that such risks are more often morally unjustifiable. But it would, in principle, be possible to have a "sliding scale" test encompassing all wrongs, and requiring stronger justification as the perceived probability of the risk approaches 100 percent.

If, instead, personal responsibility were judged *ex post*, negligence would not be a distinct topic of moral and legal inquiry. After the fact, we would simply ask whether, for example, the harm actually caused was justified in light of the benefits achieved by not taking a precaution. For example, even if you justifiably believed *ex ante* that driving your very sick child to the emergency room at a high speed was a *reasonable* risk to take, the *ex post* view asserts that if you caused property damage along the way, and if it turns out that speeding was not actually necessary to protect your child's health, then your conduct was simply *unjustified*. Conversely, even if it should have appeared to you *ex ante* that speeding to a business meeting in order to facilitate a valuable deal was an *unreasonable* risk to take, the *ex post* view asserts that if your speeding did not, as it turns out, cause anyone harm, then the speeding was *justified*.

The concept of *ex post* justification is certainly intelligible.[10] But if negligence is to be understood as an instance of personal fault, of unjustifiable risk-creation, it must be judged *ex ante*. A person's manner of driving, or a doctor's choice of medical technique, can be morally negligent whether or not it causes harm.[11] Of course, other forms of culpable behavior are also judged *ex ante*. It is culpable to attempt to kill someone, whether or not one succeeds. Negligent driving, attempted murder, and other forms of culpable conduct are culpable at least in part because of the bad consequences that the agent expects, or (if he does not expect them) that he should expect—or because of the bad consequences that he desires, or (if he does not desire them) that he should have a stronger desire to avoid.[12]

The *ex ante* analysis is not without controversy. It raises two significant problems, one normative, the other empirical. The normative problem is the question of moral luck. If I act in a particular culpable manner and a bad result follows, while you act in the same manner but a bad result does not follow, am I more to blame, because of my "bad luck"? Arguably, I am not, insofar as the difference in result might be due to factors not within my control, or factors for which I am not morally responsible: a sudden gust of wind, a third party's intervention, or the like. Those who believe

[10] A fully consequentialist account does have difficulty making sense of *ex post* justification, insofar as we cannot be certain, until the end of time, whether an act will turn out for the best. (I thank David Lyons for this point.) On the other hand, if we relativize the *ex post* judgment to the information known at the time of judgment, a qualified consequentialist assessment is possible. The idea of an *ex post* judgment of an act normally does presuppose such a relative judgment, occurring at some time subsequent to the act being judged.

[11] In law, negligence sometimes refers to unreasonable conduct, and sometimes to unreasonable conduct that incurs legal liability (usually, but not always, in the form of *ex post* compensation).

[12] The unjustifiable bringing about of a harm or death is often conceptualized as "wrongdoing" (the badness of an act), as compared to "culpability" (the blameworthiness of an actor). The latter category concerns the offender's degree of blame for bringing about a wrong, and thus includes the actor's mental states and excuses. For some doubts about this conceptualization, see Kenneth W. Simons, "Deontology, Negligence, Tort, and Crime," *Boston University Law Review* 76, nos. 1 and 2 (February/April 1996): 285–89.

to the contrary, who accept "moral luck," partially endorse the *ex post* perspective. They believe that whether the harm occurs, even if this is partly a question of luck, affects the seriousness of the wrong.

Thoughtful commentators disagree about moral luck.[13] For purposes of this essay, however, a resolution is unnecessary. For even if we do not completely reject the principle of moral luck, we should at least insist on some minimum *ex ante* responsibility as a predicate for any personal responsibility, even if we permit moral luck to *increase* blame or responsibility beyond that minimum.[14]

The empirical problem with *ex ante* analysis is the problem of hindsight bias. Cognitive scientists point out that ordinary people, and even professionals such as doctors, are "biased" in their assessment of *ex ante* probabilities in the following way: if a risky decision (such as the choice of a medical technique) is described as actually causing a harmful result, people give a much higher estimate of the *ex ante* risk than if the decision is described without reference to the result.[15] Insofar as this phenomenon is deep-seated, an unbiased, *ex ante* perspective will be very difficult to achieve. (For example, juries in tort cases would be unable to assess *ex ante* negligence fairly unless they were kept in the dark about whether the plaintiff had been harmed—a highly unrealistic option!)

A further question about the *ex ante* perspective concerns the proper understanding of *ex ante* "risk." If risk is to be an *ex ante* concept, it requires some sort of estimate of the probability of future harm (and benefit),[16] and that estimate will be based on evidence as of a certain point

[13] For endorsements of moral luck, see Thomas Nagel, "Moral Luck," in Nagel, *Mortal Questions* (Cambridge: Cambridge University Press, 1979); Bernard Williams, "Moral Luck," in Williams, *Moral Luck: Philosophical Papers, 1973–1980* (Cambridge: Cambridge University Press, 1981); and Tony Honoré, "Responsibility and Luck," *Law Quarterly Review* 104 (October 1988): 530–53. For criticism, see Feinberg, *Doing and Deserving*, 31–33; and Steven Sverdlik, "Crime and Moral Luck," *American Philosophical Quarterly* 25, no. 1 (January 1988): 79–86.

[14] See Michael S. Moore, "The Independent Moral Significance of Wrongdoing," *Journal of Contemporary Legal Issues* 5 (Spring 1994): 281; and Kenneth W. Simons, "When Is Strict Criminal Liability Just?" *Journal of Criminal Law and Criminology* 87, no. 4 (Summer 1997): 1111–12.

[15] Baruch Fischhoff, "For Those Condemned to Study the Past: Heuristics and Biases in Hindsight," in D. Kahneman, P. Slovic, and A. Tversky, eds., *Judgment under Uncertainty: Heuristics and Biases* (New York: Cambridge University Press, 1982), 341; Baruch Fischhoff, "Debiasing," in ibid., 427–31. For an experimental study finding hindsight bias in judgments of negligence, see Susan J. LaBine and Gary LaBine, "Determinations of Negligence and the Hindsight Bias," *Law and Human Behavior* 20, no. 5 (1996): 501–16.

[16] The risk that negligence analysis presupposes is typically risk about future harm, not about future benefit. Is this a necessary feature of the negligence concept? Is an actor negligent if the risk pertains only to the future *benefit* that might justify imposing the risk of harm, and not to the future harm itself? Suppose I speed my car through your rose bushes, with a high probability of causing property damage, because I believe I must bring my child to the hospital. If I am unreasonable in thinking that there is any significant health risk to my child, am I negligent in causing that damage? In a sense, I am; but the more typical sense of negligence confines the concept to low-probability risks of *harm*. This issue also arises with the use of defensive force, insofar as one might be justified even if there is only a modest probability that the use of such force will be socially beneficial (in preventing harm to the victim).

in time.[17] But should we evaluate future risk subjectively, from the perspective of the actual agent's knowledge base and capacity for inference, or instead objectively, that is, from a more ideal perspective? If we are to capture the broadest sense of negligence, the more ideal perspective is appropriate. (Otherwise, we could not describe as negligent an actor who inaccurately perceives his surroundings, or who reasons irrationally, even if his defects are due to his own culpable neglect, such as intoxication.)

Under the *ex ante* perspective, then, the concept of a negligent actor includes a person whose estimates of the risk of future harm are unreasonable or unjustifiable. And, to contrast negligence with the more serious form of culpability of *knowingly* creating a harm, I will draw a rough distinction between "low-probability" and "high-probability" estimates. The negligent actor is one who either realizes, or should realize, that she has (unjustifiably) created a "low-probability" risk of harm; while the "knowing" actor is one who either realizes, or should realize, that she has (either justifiably or unjustifiably) created a "high-probability" risk of harm.[18] As a shorthand, I will often refer to negligence as the unreasonable creation of a "risk" of harm, meaning a "low-probability" risk.

A final point about the *ex ante* perspective is as follows. The *ex ante* probability of harm distinguishes the actor who negligently creates a risk from the "knowing" actor who believes that the risk of harm is certain or almost certain. But how does the negligent actor compare with an *intentional* actor? Here we must distinguish two cases. One who intends to

[17] Probabilities can be either "objective" or "epistemic." That a coin toss will come up heads half the time is an "objective" probability; that a particular medical technique, even when carefully performed, creates a 2 percent risk of death is an "epistemic" probability. Epistemic probability is the form that is relevant to negligence. For a helpful explanation, see Stephen R. Perry, "Risk, Harm, and Responsibility," in David G. Owen, ed., *Philosophical Foundations of Tort Law* (New York: Oxford University Press, 1995), ch. 14.

One can also distinguish "risk" (probabilities that can be precisely measured) from "uncertainty" (all other probabilities). See W. Kip Viscusi, *Fatal Tradeoffs* (New York: Oxford University Press, 1992), 153–54; Nicholas Rescher, *Risk: A Philosophical Introduction to the Theory of Risk Evaluation and Management* (Washington, DC: University Press of America, 1983), ch. 8; and Clayton P. Gillette and James E. Krier, "Risk, Courts, and Agencies," *University of Pennsylvania Law Review* 138, no. 4 (April 1990): 1028 n. 1. But for purposes of understanding the general concept of negligence, this distinction is not critical.

[18] In this essay, I place quotation marks around the term "knowing" when I use the term in the special sense just noted in the text. This sense is unconventional in an important respect: we do not normally say that an agent has "knowingly" created a harm when the actor *should* have known that the harm was highly likely to occur. Rather, we reserve the term "knowingly" for one who *subjectively believes* that the harm was highly likely. I sometimes use the less conventional form in order to focus on negligence as a form of unreasonable risk-creation, and to contrast it with more risky behavior. If I were instead focusing on negligence as a form of inadvertent risk-creation, then I would contrast such inadvertence with knowledge as conventionally understood (subjective awareness of a high-probability risk of harm) and with one meaning of recklessness (subjective awareness of a low-probability risk of harm).

Negligent and "knowing" actors differ in their estimates of the *probability* of harm. A separate question is the *severity or extent* of the relevant harm. A *ceteris paribus* condition is implicit in my comparison of negligent and "knowing" actors. Negligently creating a risk of a nuclear disaster is obviously more culpable than "knowingly" stepping on someone's foot.

bring about a *harm* should not be considered merely "negligent," even if he believes that he has only a small chance of success in bringing about the harm. (The attempted or successful murderer who actually and reasonably believes that the *ex ante* chance of success is 20 percent is not merely "negligent.") On the other hand, if an agent does not intend to bring about a harm, but does intend to create a *risk* of harm, she should probably be considered merely negligent, so long as the *ex ante* probability of harm is low. (A teenager who drives near a pedestrian for the thrill of endangering him would thus be deemed negligent, but not an intentional wrongdoer.)[19]

How we should classify the intentional and the conscious creation of varying degrees of risk is inevitably arbitrary. But one point of recognizing a distinct moral and legal "negligence" category is to identify a type of culpability that is less serious, and easier to justify, than the culpability of (unjustified) knowing and intentional actors. Those who intend to cause *harm* most clearly fall within a more serious category. Those who intend neither harm nor the risk of harm clearly fall within a less serious category. The categorization of those who intend only to create a *risk* of harm is less certain; it depends on the specific account of culpability one endorses.[20]

C. What is special about negligence?

Negligence, in the sense of the (unjustifiable) creation of a future risk of harm, is generally understood to pose distinct problems and to deserve separate analysis in morality and law. But this raises a puzzle. What is so special about posing a *risk* of future harm? Does the circumstance that the actor posed a lower rather than a very high risk of harm really create a distinct moral (and legal) category, governed by distinctive principles? For the most part, I will conclude, it does not.

If negligence is not a distinctive problem, then we should not look for distinctive principles to judge whether or not a risky act is justifiable. Once we have determined the proper moral and legal analysis of harms that will occur with certainty (or of harms that have already occurred), we

[19] Of course, if she intends to cause the pedestrian fear, then she indeed intends a "harm," insofar as fear is an actual (though intangible) harm. Still, one who intends a more serious form of harm (such as physical harm) is more culpable.

It is much easier to justify intentionally creating a *risk* of a given type of harm than to justify intentionally causing that harm. As an instance of the former, consider the promoter of a trapeze act who chooses not to use a safety net, in order to make the act more exciting.

[20] On a deontological account, whether an actor intends to create a risk, as opposed to creating it as a knowing side-effect, is indeed relevant to whether, all things considered, the actor's risk-creation is justified, as I will argue below.

There is much more to say about the relevance of intention and of other conative states such as "culpable indifference" to risk or to harm. For a thorough account of the moral and legal differences between cognitive and conative mental states, see Simons, "Rethinking Mental States."

would simply apply the same analysis when the probability of harm is less. Moreover, if the widest sense of negligence is used, encompassing not only risks of *physical harm* but also the endangering of any morally relevant interest or value through human action, then negligence seems to describe nothing less than the whole of action-guiding morality.[21]

Consider, for example, the doctrine of double effect. That doctrine, supported by many deontologists, asserts that there is an important moral distinction between intending to cause a harm and knowingly bringing about a harm as a certain side-effect or further effect of what one intends: intending to harm (and thereby causing harm) is absolutely forbidden, or is subject to a heavy burden of justification, while knowingly causing harm is easier to justify. (Contrast intending to kill civilians in wartime, in order to terrorize the population, with knowingly killing the same number of civilians as a regrettable side-effect of attacking a military target.) This doctrine could also apply, however, where the probability of the harm occurring is less than a virtual certainty. Intending to create a serious *risk* of harm (and thereby causing the harm) would then be absolutely or prima facie wrong, while knowingly creating a serious *risk* of harm would be easier to justify. (Contrast Alice driving near a pedestrian for the thrill of endangering him, with Betty driving just as close to a pedestrian as an unavoidable incident of bringing her sick child to the hospital.)

The same puzzle arises with moral norms other than the doctrine of double effect. Consider the famous "trolley problem."[22] A trolley's brakes fail. If I do not turn the trolley, it will kill five workmen; if I turn the trolley onto a spur of the track, it will kill one. Should I turn the trolley, thereby causing one death; or decline to turn the trolley, with the result that five will die? It is difficult to see why the problem should change if one discounted the expected harm from each choice by an equivalent amount. Thus, suppose I know that the brakes have a 10 percent (rather than 100

[21] Ronald Milo adopts a wider version that he calls "moral negligence" to encompass "any kind of morally wrong act due to a particular kind of shortcoming on the part of the agent—namely, a culpable failure to take those precautions necessary to assure oneself, before acting, that what one proposes to do is not in violation of one's moral principles." Ronald D. Milo, *Immorality* (Princeton: Princeton University Press, 1984), 84. This is an *epistemic* duty, to ascertain whether one's act would violate moral principles; Milo points out that an additional question is whether (and in what way) we are culpable for nevertheless taking the risk of violating our principles. See ibid., 85.

[22] See Philippa Foot, "The Problem of Abortion and the Doctrine of the Double Effect," in Foot, *Virtues and Vices* (Berkeley: University of California Press, 1978), 23-24; and Judith Jarvis Thomson, "The Trolley Problem," in Thomson, *The Realm of Rights* (Cambridge: Harvard University Press, 1990), ch. 7. The trolley problem is often posed in contrast with the following "transplant" problem: a surgeon is considering whether to carve up an unwilling patient and transplant his organs as the only means available to save five other lives. See Thomson, *Realm of Rights*, 137. Some explain the impermissibility of the trade-off in the transplant case, and its permissibility in the trolley case, by reference to the doctrine of double effect (insofar as the deaths of the workmen supposedly are foreseen but not intended, while the death of the patient supposedly is intended). But I agree with Thomson and others that this explanation does not suffice.

percent) chance of failing, whichever direction the trolley takes. The expected value of the harm from turning the trolley is now 0.1 deaths (rather than one), and the expected value of the harm from *not* turning the trolley is now 0.5 deaths (rather than five). Under this variation, the same considerations still inform the moral question of whether one must, may, or may not turn the trolley. If diverting a threat under these circumstances is morally permissible when the alternative expected harms are one versus five deaths, it appears to be permissible when the alternative expected harms are a 10 percent probability of death versus a 50 percent probability (or 1 percent versus 5 percent).

Moreover, if probabilities were a critical and independent determinant of moral permissibility, then what we are permitted to do would rest on the nature of the risk description, which is sometimes arbitrary. For example, when actors engage in repetitive or far-reaching activities, it is somewhat arbitrary whether one describes the risks in temporally or spatially limited terms, or instead in more capacious terms; yet the first description will yield a lower probability than the second. The chance that I will injure someone through moderate speeding over my lifetime might be 15 percent; while the chance that my moderate speeding will injure someone in a single trip to the beach is vanishingly small.[23]

Further, even when one creates an intentional or knowing harm, one can be justified in so acting even when the reasonably foreseeable probability that the justifying facts exist is considerably *less* than one. One might, for example, be morally entitled to use defensive force so long as one reasonably believes that there is a *significant risk* that one would otherwise suffer substantial, unavoidable harm.

A number of reasons might be offered to explain why negligence is a distinctive subject of moral and legal analysis. A first reason can safely be

[23] Or, to return to the Ford Pinto example, the risk that any *individual* Ford Pinto vehicle would catch fire and cause a burn death (that reasonable precaution would have avoided) was .0000144, over a fleet of 12.5 million vehicles; but the expected number of burn deaths over the entire fleet of Ford Pintos was a probability greater than one—namely, 180 expected deaths. See Simons, "Rethinking Mental States," 292 n. 69.

To be sure, a more careful identification of the relevant frame of reference for assessing probabilities might eliminate this arbitrariness. (Compare the question of whether an individual owner must take a precaution, with the question of whether Ford must do so for all cars with the problem.) Under this approach, however, it becomes difficult to identify risks, acontextually, as "low" rather than "high" probability, i.e., as negligent rather than "knowing."

In the end, a completely acontextual identification of risks as "low" rather than "high" seems impossible. The distinction between negligence and "knowledge" appears to be a relative judgment. Consider a question that has troubled legal scholars: whether a manufacturer of a product who knows that a small number of users (out of a much larger class) are virtually certain to suffer physical harm should be treated as "knowingly" inflicting that harm. If the issue is whether his conduct demands as strong a justification as a manufacturer who knows that every user will suffer the same degree of physical harm, the answer is clearly no. And if the issue is whether the conduct of either manufacturer would demand as strong a justification if the risks of harm were substantially lower, again the answer is clearly no. But there might be no nonarbitrary way to characterize any of these four cases as "negligent" or "knowing" in an absolute sense.

put to one side. This is the point that many negligent acts involve a risk that will not eventuate in physical harm for a considerable period of time. Thus, such acts might, in the interim, cause significant emotional harm to potential victims. (Consider the fear that the presence of toxic chemicals induces in exposed populations.) But this point does not distinguish negligent from higher-probability harms; for it depends on latency, not on the fact that the risk is less than certain.[24]

A second, and more persuasive, reason why negligence is considered a distinctive moral and legal concept flows from the following fact: people often react to *lower*-probability risks of future harm in a distinctive way. For example, people often overestimate the magnitude of very small probabilities.[25] Moreover, when probabilities are lower, people might be more likely to differ among themselves, and more likely to differ with "expert" assessments, about the magnitude of risk. These differences could be due to such cognitive heuristics as framing (whether an option is characterized as suffering a loss, or instead as failing to obtain a possible benefit), the availability heuristic (whether similar instances easily come to mind), anchoring (where people have difficulty altering their initial estimates), and representativeness (whether others report high or low risks of the same phenomenon).[26] But these variations in perspectives are likely to be much less when the expected probabilities are high (for example, if a new food product is almost certain to cause slight indigestion in all users). One plausible response to these variations in perspectives is the effort to develop a unifying methodology; and a utilitarian calculus based on subjective preferences is an obvious candidate. In short, the permissibility of imposing lower-probability risks might be thought to demand a distinctive *type* of moral justification.[27]

A third reason for distinctive treatment only warrants a quantitative, not a qualitative, distinction between negligence, on the one hand, and higher-probability or "knowing" harms, on the other. This is the point that, often, the smaller the risk of harm, the easier it is to justify creating that risk. Thus, intentionally causing the death of another human being

[24] Thus, one might know to a certainty that one will suffer a harm either in the immediate *or* in the distant future; the distant harm might cause a different type or degree of emotional stress than the immediate harm. To be sure, the contemporary emotional harm produced by long-latency risks that are virtually *certain* to result in ultimate harm will often be disproportionately greater than the harm produced by less certain risks. (It will often be more than five times as painful to worry about a virtually certain future death than to worry about a 20 percent risk of death.) But these emotional harm cases do not justify distinct treatment of negligence in general. Not all cases of risk generate significant emotional harm. (In many cases, the risk is unknown or underappreciated.) At most, the considerations just discussed would justify special treatment of certain emotional harm cases.

[25] See Viscusi, *Fatal Tradeoffs*, 150.

[26] See Gillette and Krier, "Risk, Courts, and Agencies," 1091-93; and Richard H. Pildes and Cass R. Sunstein, "Reinventing the Regulatory State," *University of Chicago Law Review* 62, no. 1 (Winter 1995): 55-64.

[27] In the end, however, I will reject this argument.

when one believes that that result is a certainty is almost always (or, on some views, always) unjustifiable; but intentionally creating what one believes to be a modest *risk* of death is more often justifiable. Deliberately running someone over with your car is obviously more difficult to justify than deliberately creating a risk that you will run someone over.

A fourth reason, related to the previous one, is the possibility that the deontological constraint against knowingly or intentionally causing a person serious harm is *much* more stringent than the constraint against negligently causing such a harm. The difference might be qualitative, not merely quantitative. Consider the following two examples (for which I thank Leo Katz). Suppose Alfa speeds in the vicinity of a pedestrian, knowing that she is almost certain to kill him, because this is the only way to save the lives of five passengers whom she must bring to the hospital. Benna speeds in the vicinity of a pedestrian, aware that she is running a 20 percent risk of killing him, because that is the only way to save the life of *one* passenger whom she must bring to the hospital. Many would conclude that Alfa acted impermissibly while Benna acted permissibly. Most would at least conclude that justifying Alfa's conduct is more difficult than justifying Benna's. And yet the justifying benefits in each case are five times the expected harm. (Put another way, if 20 Alfa-situations and 100 Benna-situations arise each year, then the Alfa-situations and Benna-situations will each result in 20 deaths and the saving of 100 lives annually.)

These examples suggest a special moral concern, and a constraint of special stringency, when an actor creates a very high risk of killing another. The concern is not just an extension of the principles of justification that apply when one creates a much lower risk of death. For the constraint requires *more* than a simple proportional increase in the justifying benefit to correspond to a similar increase in the level of risk.[28]

With this important caveat, I conclude that negligence is not as distinctive a topic of moral and legal inquiry as it is often believed to be. Negligence often does not demand special analysis, except as a matter of degree. For example, the basic point of the doctrine of double effect, that knowingly causing X is much easier to justify than intending to cause X, is no less valid when X is a small risk of harm than when it is a virtual certainty of harm. (Consider a variation of the terror-bombing example: intentionally exposing a group of noncombatants even to a small *risk* of future harm is more difficult to justify than exposing them to such a risk

[28] To some extent, the law recognizes this distinction, for it sometimes shifts the burden of persuasion, and narrows the grounds of justification, when the conduct moves from negligent to knowing or intentional. See note 9 *supra*.

I have suggested that a constraint of special stringency applies when a person knowingly or intentionally creates a high risk of *death*. Whether an unusually stringent constraint also applies to a person who knowingly or intentionally creates a high risk of a *lesser* harm than death is less certain, but I cannot explore the matter here.

as a known but regrettable side-effect of a purely military action.) Thus, normative principles that apply to risks one reasonably believes are certain to occur should ordinarily also apply to lower-probability risks, *mutatis mutandis*. (On the other hand, as a *factual* matter, the creation of lower-probability risks will more often be justifiable.)

III. Private Morality

This section examines negligence, not as a question of law, but as an issue of personal responsibility. How should we choose between risk and precaution?

Although most moral theorists have given far more attention to intentional and knowing harms, we should consider how different moral perspectives would affect the analysis of negligence. In this section, my approach is to examine common-sense moral judgments and convictions about risk, and also to relate these to more general consequentialist and deontological frameworks.

Utilitarianism is often offered as the best account, or even the only plausible account, of when risky conduct is justifiable.[29] I thus begin with this perspective, before examining others.

A. Utilitarianism

Some would analyze negligence as a straightforward question of rational choice, as follows. If an agent would incur all the costs and reap all the benefits of risky action, she would rationally maximize the benefits and minimize the costs. Then we could extrapolate from this intrapersonal case to the interpersonal case. This approach has some initial plausibility.

Suppose you are a hermit living in the woods, and you are trying to decide how sturdy a deck to build as an addition to your house. Simplifying, you might consider three options—a flimsy deck, a sturdy deck, and a super-sturdy deck. What would be relevant to your decision? The sturdier the deck, the more costly it will be, in terms of labor and materials. But a stronger deck will last longer, and it will also be safer. It might also offer other advantages, including the ability to hold grills and lawn furniture, or firewood, or your private sculpture collection.

Most people would consider factors such as these in deciding what type of deck to build. Many of the factors require at least rough estimates of probability. How long will a "sturdy" deck last? We know that the probability of its lasting one year is very high; ten years is fairly high; one hundred years, perhaps low. Similarly, how likely is an injury? How do you expect to use the deck—for eating? Reading? Training your pet lion? Moreover, most of us would (if only implicitly) normally consider the *marginal* costs

[29] See Heidi M. Hurd, "The Deontology of Negligence," *Boston University Law Review* 76, nos. 1 and 2 (February/April 1996): 249–72.

and benefits. Is it worth the extra money to build a super-sturdy deck, as opposed to a sturdy deck, relative to the extra durability, safety, and functionality benefits (discounted according to their probabilities)?

This intrapersonal example provides intuitive support for a norm of rational egoism. It seems plainly irrational not to consider all the costs and benefits of alternatives. The prudent course, it appears, is to maximize benefits and minimize costs. It is then tempting to extend the intrapersonal analysis to the interpersonal context.

Suppose you, the hermit, decide to rejoin society. You expect to invite others to your house. Now the expected benefits of building a deck include the benefits to others, the sociability benefits to you, and the collective benefits of friendship and community. And the expected costs might include greater expenditures of labor and material, and new risks of personal injury to others.

Most of these new costs and benefits are similar in kind to the costs and benefits in the intrapersonal case. Why not conclude, then, that the maximizing approach appropriate in the intrapersonal case is also appropriate in the interpersonal case?

Some have so concluded. For example, in his influential book *Economic Analysis of Law*, Richard Posner employs precisely this argument.[30] An individual will balance the marginal cost or burden (B) of taking a precaution against the marginal probability (P) and magnitude (L) of the loss to that individual if the precaution is not taken. If B is less than P times L, the individual will rationally take the precaution.[31] But if the losses are to others, we need a legal liability rule to ensure that the individual takes the correct precaution. (The idea that courts should balance B, P, and L to determine whether an actor is negligent was espoused by Judge Learned Hand in the now-famous "Learned Hand formula.")[32]

This is the informal, intuitive case for a maximizing, aggregative conception of negligence. Reasonable care in the choice of risky activities

[30] Richard A. Posner, *Economic Analysis of Law*, 5th ed. (New York: Aspen Law and Business, 1998), 179–83.

[31] Actually, this test (known as the "BPL" test) contemplates that the rational actor would aggregate the different risks that would be prevented by a precaution, e.g., risks of minor physical injury (discounted by their probability), of major physical injury (also discounted), of death, of major property damage, of minor property damage, and so forth.

[32] The landmark case is *United States v. Carroll Towing Co.*, 159 F.2d 169, 173 (2d Cir. 1947). Whether Judge Learned Hand actually intended his "BPL" test to be the sort of cost-benefit economic test that Posner defends is a controversial question.

In his early writing, Richard Epstein analyzes the contrast differently, concluding that extrapolation from the intrapersonal to the interpersonal case justifies a general rule of *strict* liability. However, his focus is not on which decision among risky alternatives is best, but rather on who should bear the costs of the decision. In the intrapersonal case, he points out, all the costs and benefits accrue to the actor. In the interpersonal case, however, the actor might derive the benefits while the victim might bear the costs. Epstein suggests that the actor has no right to dump the costs of his action on another. Richard A. Epstein, "A Theory of Strict Liability," *Journal of Legal Studies* 2, no. 1 (January 1973): 159. For a critique of this argument, see Ernest J. Weinrib, *The Idea of Private Law* (Cambridge: Harvard University Press, 1995), 171–75.

consists of maximizing the net differential between benefits and costs. What is wrong with this conception?

The problems are many. First, even in the intrapersonal case, it is not true that an actor ought always to maximize the net benefit. A maximizing strategy, because of the *way* in which it values, would undermine some kinds of morally important relationships and interests.[33] For example, it is not wrong to weigh financial gains against the value of friendship in some circumstances (as when one is deciding whether to move to a new city in order to obtain a more lucrative job). But it does not follow that friendship always has a price attached to it; one should not break a date with a friend because a third person offers payment to do so. The latter practice would undermine the value of friendship in a way that the former would not.[34] More pertinent to a discussion of negligence, one should not create what would otherwise be reasonable risks to oneself in order to promote a personal value that would be inconsistent with one's own objective self-worth. For example, one should not drive at very high speed on a deserted street simply to experience the thrill of a near-death experience; but one may, and perhaps should, drive at an equivalent speed in order to reach a hospital in time to save one's own life.

Second, the intrapersonal case normally raises no problem about *consent*. By contrast, in the interpersonal case, others endangered by your risky conduct might not consent at all; or if, in some sense, they do accept the risk, they often do not consent in as full a sense as you do in the intrapersonal case.

Third, the extrapolation to the interpersonal case implicitly characterizes society as a kind of interest-maximizing "super-person." Yet, as many critics of utilitarianism have argued, this characterization ignores the point that society is composed of individuals, each with his or her own life to lead.[35] A principle that aggregates the welfare or utility of all persons (and then requires maximization of that utility) thus needs a distinctive justification.[36]

Fourth, the extrapolation is indeterminate and potentially both too weak and too strong. The constraints on risky action are too *weak* if the original actor (in the intrapersonal case) is one who took very little interest in his own safety; for the result of the interpersonal extrapolation is that he will be justified in imposing enormous risks of harm on others relative to small benefits to himself. (At the same time, however, he will be willing

[33] See, generally, Elizabeth Anderson, *Value in Ethics and Economics* (Cambridge: Harvard University Press, 1993), 66–73.

[34] See ibid., 70, discussing Joseph Raz, *The Morality of Freedom* (Oxford: Clarendon Press, 1986), 349.

[35] See John Rawls, *A Theory of Justice* (Cambridge: Harvard University Press, 1971), 27–29; see also Bernard Williams, in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (Cambridge: Cambridge University Press, 1973), 116–17; and Raz, *Morality of Freedom*, 271–87.

[36] At the same time, simple extrapolation will *fail* to capture the collective benefits that can be achieved only through social interaction. When and only when the hermit rejoins society, the collective benefits of friendship and community are possible.

to *accept* enormous risks to himself relative to small benefits to others.) On the other hand, the constraints on risky action might be considered too *strong* if the original actor is one who placed extraordinary value on his own safety. For then the interpersonal result is that he will only be justified in imposing slight risks to others, even when the benefits to himself are great. (At the same time, he will be willing to accept only small risks to himself even when the benefits to others are great.)

Can these problems be overcome by a less catholic view of what it is reasonable to do in the intrapersonal case? Perhaps a reasonable regard for your own safety does require that you give it at least weight X (or precisely weight Y). Then we could extrapolate that you should give the safety of others at least weight X (or precisely weight Y) in the interpersonal case.

Still, this extrapolation is problematic. In the intrapersonal case, it is highly implausible that all persons must give any precise weight (Y) to any of their ethically relevant interests, including their safety. At most, you might be ethically obliged to give your safety at least a certain weight (X), while you would be permitted but not required to give it more weight. But even if this is so, there is still no reason to assume that you must give weight X, but need not give more than weight X, to the safety interests of others when your actions affect them. It is quite plausible that you are entitled to show very little regard for your own safety or health relative to your other interests (such as love of excitement, desire not to burden others, and the like). It is not plausible that you can extrapolate this entitlement to the situation in which the safety or health risks are imposed on others while the corresponding benefits belong to you.[37]

Another way to see the problem with extrapolation is to compare a case in which a third party is in a position where she must make a decision for

[37] In his attempt to explain American tort law, Stephen Gilles employs the extrapolation approach (which he terms the "single-owner" heuristic) in an especially interesting way. He asks what value the average injurer would assign to precaution costs. But, recognizing that the average injurer might assign too low a value to the expected accident costs to others, Gilles also asks what value the average *victim* would assign to those costs. "Because the average injurer and the average victim, taken together, constitute the average *person*, the inquiry reduces to whether the average person would take the precaution if he or she bore both the costs and benefits in full." Stephen G. Gilles, "The Invisible Hand Formula," *Virginia Law Review* 80, no. 5 (August 1994): 1035.

In American tort law, however, the negligence test employs a "reasonable" or "ideal" valuation, not an "average" valuation. Gilles tries to handle this objection by referring to the accident valuation of victims, not of injurers. Still, it seems that he should refer to the "reasonable" precaution valuation of injurers. For the average injurer might place undue weight on the cost of certain precautions (e.g., the average Boston driver is probably unduly worried about his pride when he refuses to allow other drivers to share the road).

Gilles ultimately settles on an "altruistic reasonable person standard," asking what care an "average reasonable person takes of his or her own person and property" (ibid., 1037, 1038). The problem remains, however, that "average" and "reasonable" (or ideal) standards can deviate.

If one moves from an "average" to a "reasonable" valuation, one has moved from a factual, descriptive account of utilities to a normative, social valuation. (See the discussion below of a modified utilitarian calculus.)

the benefit of one person with a case in which she must make a decision for the benefit of two or more persons. Consider the following argument from Judith Jarvis Thomson (with her example slightly altered): If D is in a position where she must decide what is best for Y, she is permitted to make some relatively close marginal trade-offs (for example, cutting out unconscious Y's kidney to save Y's life) which we would not countenance if D were in a position where she must decide what is best for both X *and* Y (cutting out unconscious X's kidney to save Y). To be permissible, the trade-off in the latter case must tilt much more strongly in Y's favor than the trade-off in the former case.[38]

In the end, extrapolation is unpersuasive. People differ greatly in their concern about their own safety, not to mention the type and strength of their other interests, purposes, and values that would (intrapersonally) justify risking their own safety. (Some place enormous value on their own health, others on the pleasures of risk-taking, or on saving time or expense; and so forth.) People have a moral prerogative or permission, within a rather wide range, to balance these interests in many different ways. Yet it is doubtful that moral principles for *interpersonal* risk-imposition should be so variable or that the acceptable range should be so wide. The interpersonal variations are complex and often unknown to the risk-imposer (or victim). Just as important, accommodating these differences within an interpersonal moral norm governing risk-imposition seems wrong in principle, not just difficult in practice. I should not feel free to ride my bicycle more quickly and recklessly in the vicinity of a depressed or masochistic person than in the vicinity of a person who attaches extraordinary value to his personal appearance and thus is extraordinarily averse to personal injury.

On the other hand, the utilitarian approach does have some attractive features. It demands that we consider carefully all the consequences of our actions. In the particular context of negligence, this demand seems especially apt, since even the immediate consequences are variable and, by definition, not highly probable. Choosing a particular risky course of action A as opposed to course of action B (or no action C) might have numerous possible consequences, differing in their probability and their magnitude if they occur. Utilitarianism offers a method for combining these disparate consequences via a single formula.

B. Modified utilitarianism

Utilitarians are not without responses to some of the above problems. In several ways, they might tinker with the utility calculus to bring it

[38] See Thomson, *Realm of Rights*, 197–99. It might be permissible to take *blood* from unconscious X in order to save Y.

closer to ordinary judgments about risk and closer to common-sense morality.

First, consider ordinary attitudes toward risk. Clayton Gillette and James Krier point out that lay opinions about risk, although they systematically differ from expert opinions in various ways, are not necessarily inferior. Experts tend to focus only on overall mortality or morbidity in comparing risks: "a death is a death." Lay opinions tend to be much richer and more nuanced. This richness might reveal, not irrationality, but different values—for example, justifiable concerns about whether a risk has been voluntarily incurred, or will result in a catastrophic harm, or will have delayed effects, or is irreversible in its effects, or is man-made as opposed to natural.[39] Still, Gillette and Krier seem to assume that these values can be accommodated within a broadly utilitarian, "cost-minimization" framework.[40] Extending their analysis, one might add a "premium" to those risks that are involuntarily as opposed to voluntarily incurred. And one might similarly modify the utilitarian calculus to accommodate the cognitive heuristics (such as framing and anchoring, mentioned above) that laypersons often use.

Second, utilitarianism is often criticized for its indifference to how benefits and burdens are distributed. But a broader form of consequentialism can accommodate concerns about the fairness of distribution.[41] For example, the best action might be that which maximizes benefits, subject to a distributive condition. The condition might be ensuring a minimum level of benefits for all affected persons, or not producing certain extremes of inequality (especially with respect to imposing losses rather than failing to confer benefits), or not concentrating large harms (or the risk of large harms) on certain individuals.

We must be careful, however, to distinguish two questions: (1) whether the fairness of distribution is relevant to the action a person may or should take, and (2) how the costs and benefits of a (concededly) permissible action should be distributed. Negligence as a form of culpability or wrongdoing addresses only the first problem. (In tort law, strict liability,

[39] Gillette and Krier, "Risk, Courts, and Agencies," 1071–79.

[40] Ibid., 1028 n. 2. They do, however, acknowledge a possible role for distributive principles.

[41] See, e.g., Derek Parfit, "Equality or Priority?" The Lindley Lecture (University of Kansas, November 21, 1991); and David O. Brink, *Moral Realism and the Foundations of Ethics* (New York: Cambridge University Press, 1989), 270–73 (where Brink argues that an objective form of utilitarianism can endorse a distribution-sensitive theory of value).

Concern about distributive effects is a powerful reason not to adopt Richard Posner's suggested wealth-maximization version of utilitarianism. This version evaluates choices by the criteria of willingness and ability to pay, rather than utility. But these criteria create an additional problem of distributive justice, beyond that entailed by utilitarianism. For example, under the wealth-maximization approach, it is better to endanger the safety of a poor person than the safety of a wealthy person: "a person should feel free to drive faster in a poor than in a wealthy neighborhood because expected accident costs are on average lower in the former," as Posner candidly concedes. Richard A. Posner, "Wealth Maximization and Tort Law: A Philosophical Inquiry," in Owen, ed., *Philosophical Foundations*, 110.

or liability without fault, addresses the second problem: even though the strictly liable actor might have acted permissibly, fairness or some other principle requires him to pay the costs of his action.)[42]

Often, in deciding what action a person is permitted or required to take, we properly do *not* consider how the costs and benefits ultimately should be distributed. The classic examples are necessity cases: a starving backpacker is permitted to break into an unoccupied cabin and steal some food, and a ship owner is permitted to damage another's dock to save his crew or even his ship, even though each should ultimately pay for the goods that he consumes or damages.[43]

Sometimes, however, distribution is relevant to what an agent may or should do in the first instance. Suppose a thief issues this demand: "I'll take $500 from X, or instead I'll take $10 from each of seventy people. You choose." A proper concern for fair distribution of losses supports your choosing the second option. Or, more relevant to the negligence debate, suppose a company is deciding how to dispose of toxic waste, either by imposing a modest risk of future harm on a community that is poor, or that has already been the dumping ground for waste; or by imposing a slightly higher risk of harm on a more wealthy or more pristine community. The latter choice is certainly defensible. One's choice of location should be sensitive to the distribution of the risk of harm, as well as to its aggregate amount; and in some cases, it would be better to distribute a larger amount of risk if that is the only way to distribute the risk more fairly.[44]

[42] Criminal law does not redistribute costs. Therefore, if strict criminal liability is justifiable, the justification must be different. See Jules Coleman, *Risks and Wrongs* (New York: Cambridge University Press, 1992), 222-23; and Simons, "When Is Strict Criminal Liability Just?"

[43] However, it is also proper to consider not only the actor's primary conduct, but also his ability to insure against the risks of his conduct. The permissibility of engaging in some activity might itself depend on ability to absorb certain risks of one's conduct. We forbid people from driving without insurance, in part because we want them to be financially responsible even for the non-negligent accidents they cause. See Kenneth W. Simons, "Jules Coleman and Corrective Justice in Tort Law: A Critique and Reformulation," *Harvard Journal of Law and Public Policy* 15, no. 3 (Summer 1992): 880. See also Thomson, *Realm of Rights*, 159.

[44] What counts as a fair distribution is beyond the scope of this essay. In environmental law, government regulators have increasingly attended to the distribution of risk as well as its aggregate level. See Pildes and Sunstein, "Reinventing the Regulatory State," 44.

The special concern that many feel about "catastrophic" loss is partly based on a distributive concern:

> Imagine . . . a decision maker who is forced to choose between two actions. The first action poses a 1 in 1,000 chance of causing 100,000 deaths spread randomly across the country; the second has a 1 in 1,200 chance of causing the near obliteration of a city of 100,000. A rational decision maker could obviously select the first alternative, notwithstanding its larger expected loss.

Gillette and Krier, "Risk, Courts, and Agencies," 1078.

Moreover, imposing a large harm (or the risk of a large harm) on a single person is often impermissible when it is offset only by small benefits to each of a large number of persons,[45] and even when it is offset only by the avoidance of small *losses* to each of a large number of persons. As an instance of the latter, consider an example from Ronald Dworkin: "Suppose ... that my child's life depends on a noisy ambulance that annoys a large number of people who would collectively pay more not to be annoyed than all the funds I have."[46]

A third way in which utilitarianism can be modified to accommodate some of the above criticisms is that preferences can be filtered or "laundered."[47] A more objective form of utilitarianism can replace private preferences, pleasure, happiness, or utility, with "social" utility.[48] On this view, for example, the utility calculus would give no weight to the pleasure that a reckless driver obtains from the thrill of endangering others, or to the pleasure that racists obtain from the knowledge that they have created an environmental hazard that disproportionately endangers blacks.

Of course, this qualification of utility is itself problematical in several ways. The qualification undermines some of the advantages of the utilitarian approach, including its neutrality among preferences or forms of utility. It is also indeterminate. How does one distinguish items with no (or even negative) social utility from those with positive social utility? If we give no weight to a person's thrill from endangering others, should we similarly give no weight to a person's thrill from driving at high speed, where the thrill does not actually depend on the risk to others' safety?

Moreover, the laundering approach can be ad hoc; sometimes it appears to reflect *non*utilitarian moral judgments, not a neutral criterion of "socially acceptable" utility. And, most importantly, the approach sometimes misdescribes the *way* in which the relevant moral principle operates. For example, many utilitarians would grant that racist or sadistic preferences should not be even part of the justification of actions. But do we best reflect this moral principle by modifying the general utilitarian calculus to ignore such preferences? Consider a case in which the actual motive of the actor was racist or sadistic, but other legitimate justifications were available to him. (Suppose an employer fires a minority employee for racist reasons, but, unknown to the employer, the employee

[45] See Thomson, *Realm of Rights*, 166–68; and Raz, *Morality of Freedom*, 276. For example, the conclusion that it is wrong to humiliate another for fun is unaffected by the number of persons who would derive pleasure from such an act. Anderson, *Value in Ethics and Economics*, 69.

[46] Ronald Dworkin, *Law's Empire* (Cambridge: Harvard University Press, 1986), 307.

[47] See Robert E. Goodin, *Utilitarianism as a Public Philosophy* (Cambridge: Cambridge University Press, 1995), ch. 9 ("Laundering Preferences").

[48] See Brink, *Moral Realism*.

was embezzling funds.) A properly "laundered" utilitarian calculus might still consider the action justifiable, while a nonutilitarian approach might condemn the action, notwithstanding its utilitarian net benefits.[49]

This concern that "laundering" mischaracterizes the underlying moral principles applies as well to the first modification of utilitarianism discussed above—namely, translating the special characterization of certain risks (as "involuntary" rather than "voluntary," or man-made rather than natural) into a quantitative "premium" in the utilitarian calculus. Again, we might not correctly capture the reason why we are concerned about "involuntary" risk impositions if we were simply to apply a risk premium (for example, if we first computed the appropriate value of a "voluntary" risk, and then, in the case of an "involuntary" risk, multiplied the first value by two). Rather, fully voluntary acceptance of certain risks (such as the risks of experimental medical treatment) might mean that the creation of those risks is not wrongful at all. And, for risks that vary in their "involuntariness," the best analysis might have a different structure than a utilitarian calculus. (It might, for example, forbid the imposition of any high-level involuntary risks, and permit the imposition of low-level involuntary risks only if the risks either are widely shared in the community or are absorbed only for a limited period of time.)

C. *Nonutilitarian approaches*

1. Rejecting the utilitarian rationale itself. A more basic critique rejects the utilitarian notion that maximizing the aggregate good consequences for all persons is the correct moral principle for problems of risky action. First, it appears that the utilitarian approach cannot definitively characterize negligent acts as instances of wrongdoing. Second, the utilitarian approach can appear heartless and cold-blooded. The first critique, I will suggest, is more potent than the second.

A comprehensive utilitarian approach considers all consequences that affect utility. The consequences include the long-term as well as the short-term, and the best decision procedure as well as the best decision (considered in the abstract). On such an approach, the right thing to do in balancing risks of harm against benefits does not merely depend on which choice has the greater expected utility as measured by *immediate* consequences (such as anticipated risks of injury to others, and anticipated benefits to the actor).

For example, in deciding how fast to ride my bicycle, I should not consider only whether the immediate risks of injury to others outweigh

[49] In Goodin's terminology, the prohibition on racially motivated actions might reflect a violation of rights and thus might be better understood as an "output filter," not the "input filter" accomplished by laundering preferences. Goodin, *Utilitarianism*, 133–37.

the immediate benefits to myself and others (in terms of pleasure, speedily arriving at my destination, and so forth). In principle, I should also consider whether my decision will cause others (such as pedestrians) to adjust their behavior in specific ways, and, more subtly, whether my decision will reinforce or undermine social norms governing risky behavior.

Moreover, we should also consider, as a second-order question, which individual decision procedures should be employed. Should agents, in the process of deciding what to do, explicitly try to maximize benefits and minimize costs? Should they instead defer to certain rules of thumb or norms about safe behavior (such as conforming to social custom)? Should they be entirely benevolent, ignoring all of their own interests? Should they be entirely egoistic, considering only their own interests? Should they delegate decision making, where feasible, to a parent, or a trusted friend? In principle, any of these approaches might turn out to be optimal (considering both the immediate costs and benefits of the action and the costs and benefits of varying decision-making procedures for producing the optimal action). More likely, different approaches will be optimal in different circumstances (such as consumption of goods in a market, or health decisions when one is incompetent).

A utilitarian approach can indeed address these complications, but in doing so, it will lose much of its supposed simplicity. More fundamentally, however, utilitarianism encounters great difficulty in accounting for the judgment that negligent acts are a species of wrongdoing or culpability. For utilitarianism (and consequentialism generally) cannot fully explain the retrospective orientation of important categories of moral judgment. Suppose Emily rides her bicycle too fast, in light of the expected risks and benefits, and thereby endangers you. She is negligent if, from the *ex ante* perspective, her balancing of the risks and benefits was unreasonable (whether or not her conduct results in harm). Having done wrong, she deserves censure, and she might incur other duties: perhaps she ought to apologize, or take affirmative steps to give you medical or other aid, or pay for the resulting harm. Why is she subject to blame or to these other duties? On the utilitarian account, it is because holding her to be under these duties will produce the best overall consequences. That means we must wait, as it were, to see how things turn out, before we can confidently blame her (or before we can justify any other duty). Will blaming her have desirable consequences overall? It might, but it might not. Blaming her might be misunderstood by her devoted friends and family, blinded by love, as an *ad hominem*, undeserved attack; at the same time, blaming her might only have the most trivial effect in strengthening the social norm against negligent bicycling, or in discouraging her own negligent bicycling.

One of the most important nonconsequentialist features of retributive-justice accounts of criminal law, and of corrective-justice accounts of tort

law, is their retrospective orientation.[50] In deciding whether to punish a negligent act, or to enjoin it, or to assess damages for the harm that results from it, we need not wait to see how many future acts of negligence these responses would deter. And the same is true of our simple decision whether to blame the actor as wrong or culpable. (Nor is it the case that we must at least be able to predict that such future deterrent benefits are reasonably likely.)

Of course, this retrospective orientation does, in the case of negligence, coexist with a prospective and consequentialist feature—the feature that the action is culpable or wrong because it unreasonably risks a *future* harm, a harm that would be a consequence of the action. But this consequentialist feature is limited, and does not undermine the essentially nonconsequentialist quality of blaming, of retributive justice, or of corrective justice.[51]

Another problem with utilitarianism is that, in some versions, it turns subtle and important questions of personal or social choice into problems of calculation. Permitting these difficult trade-offs to depend on the result of a computation seems to devalue the interests at stake. It can even seem cold-blooded.

This concern is part of the reason why the Ford Pinto example troubles so many observers. A consequentialist has some powerful replies, however. First, he might say that the objection ignores reality. Our choices do have consequences, which we should not ignore. Indeed, whatever choice we make might have unfortunate consequences for human welfare: the choice against precaution increases risks of personal injury, but the choice to take a precaution might interfere with the function of a product, or price it out of reach of a significant number of consumers. One who claims that we should never trade off life or bodily injury against other interests is asserting a moral position that is unrealistic or simply fanatical.

Second, one of the reasons that the Ford Pinto example provokes outrage is a sense that the actor, having made a justifiable choice, might feel free to ignore the negative consequences of his decision. But, the consequentialist could point out, this conclusion need not follow. One might defend an actor's making a calculated trade-off, but at the same time insist that he has secondary duties with respect to the harm that he

[50] Strictly speaking, the orientation is "non-prospective" rather than retrospective. Our judgment that a negligent act deserves retributive blame could be contemporaneous with (or even prior to) the act. What I want to distinguish (as "prospective") is a judgment that depends on whether the *further* consequences of the act are optimal.

[51] Indeed, the *ex ante* and action-guiding features of the negligence perspective mean that negligence is *necessarily* a consequentialist doctrine. But this is so only in the very limited sense that the harmful consequences immediately risked by the negligent act are critical to the actor's culpability. In precisely the same sense, the wrongfulness of attempted murder, too, depends on the expected or intended consequences of the attempted murderer's acts. Only in this limited sense is our reason for blaming negligent actors or attempted murderers necessarily "consequentialist."

justifiably causes—including a duty to inform the potential victims, to apologize, and sometimes even to compensate for some or all of the harm that they will suffer.[52]

2. *Nonutilitarian balancing.* How can one endorse the careful evaluation of (at least some) consequences without turning the problem of choice into a problem of calculation? First, one should concede that there is normally no single metric—such as money, or wealth, or utility (in any nontrivial sense)—into which all relevant values can be translated. Commensurability in this strong sense is normally unattainable.[53] Second, at a minimum the moral agent could explicitly consider a list of relevant factors, ignoring some factors as irrelevant, and identifying the direction in which the relevant factors point.[54] The Learned Hand test (discussed earlier) could itself be so understood. It instructs that the following factor militates against taking the risk and *in favor* of taking a precaution: namely, the expected accident costs if no precaution is taken. And it might instruct us to count the following factors as militating *against* taking a precaution: the additional risks that taking a precaution would create, and the *socially reasonable* costs to the agent.[55] (This last qualification excludes such "costs" as loss of pleasure from indulging sadistic and racist preferences, as discussed above.)

This approach does not yet provide much guidance or much predictability. But can we really do any better? On one view, moral judgment consists only in an intuitionistic (that is, pluralistic and essentially unstructured) balancing of reasons.[56]

A nonutilitarian account *can* give a more principled explanation than this of how we should balance or trade off values. If we balance, even at the margin, the analysis need not collapse into a maximizing, utilitarian framework.

Balancing need not be utilitarian, or even consequentialist. For example, in deciding whether to turn a trolley and divert a threat from five people onto one, or in deciding whether to stop the trolley by throwing

[52] We might also ask producers of dangerous products to contribute to a social insurance fund for the victims.

On the general idea that tragic choices require a morally sensitive actor to show regret, see Martha Nussbaum, *The Fragility of Goodness: Luck and Ethics in Greek Tragedy and Philosophy* (Cambridge: Cambridge University Press, 1986). See, generally, Jeremy Waldron, "Rights in Conflict," in Waldron, *Liberal Rights* (Cambridge: Cambridge University Press, 1993), 214–15.

[53] See, generally, Cass Sunstein, "Incommensurability and Valuation in Law," *Michigan Law Review* 92, no. 4 (February 1994): 779–861.

[54] For an application of this approach to social regulation of risk, see Pildes and Sunstein, "Reinventing the Regulatory State," 64–66, 127 (where they endorse the "disaggregation" of costs and benefits so that citizens and decision makers make value choices more openly).

[55] In American tort law, the *Restatement (Second) of Torts* catalogues relevant *social* interests to be balanced. And the commentary to the *Restatement* does not emphasize the incentive effects of the balancing test that the *Restatement* endorses, as one would expect if the test were thoroughly consequentialist. See *Restatement (Second) of Torts*, sections 291–93 (1965).

[56] See Rawls, *A Theory of Justice*, 34–40 (where Rawls criticizes such intuitionism, and also distinguishes it from metaethical and epistemological intuitionism).

a passenger off the trolley onto the tracks, I do need to make comparisons: but I might be attending only to the stringency or relative weight of the possible victims' claims against me, or of my duties not to kill them, and not to the long-term *consequences* of permitting this type of action.[57] In the context of risky action, it is just not plausible that one has an absolute, exceptionless duty not to create any risks greater than some specified level;[58] nor is it the case that if my duty is not absolute, the limits to the duty must be consequentialist.[59] Or, to take a more prosaic example, my duty to tell the truth can conflict with my duty of loyalty to a friend. In deciding whether to tell a benevolent lie to protect my friend's feelings, I might weigh the stringency of each duty, and the extent to which alternative actions would undermine the value of friendship and the values underlying a duty to tell the truth. But such a balance is consistent with many deontological accounts, and is quite different from a consequentialist summing up of the net future benefits and harms flowing from my action.

On the other hand, consequentialism, and in particular utilitarianism, appears to offer a relatively simple, and therefore attractive, method of balancing. It is thus understandable that balancing and trade-offs (especially at the margin) are closely identified with utilitarian analysis.[60]

Is it possible to develop a clear, nonconsequentialist formula for negligence, one that accommodates competing values but avoids the problems of a pure (or even a distribution-sensitive) maximizing approach? Consider two efforts—a "disproportion" test, and a "freedom versus security" balancing test. I will conclude that these efforts, while promising, are inadequate. The first is too ill-defined, while the second is too reductionist to capture the full array of values that should be balanced.

a. Disproportion test. One possibility is a disproportion test. On this approach, if an injurer's risky conduct would expose potential victims to expected risks of $P \times L$ (probability of loss times magnitude of loss) and could be avoided only at marginal cost B, then, in order for the injurer to

[57] See Thomson, *Realm of Rights*, chs. 6 (on trade-offs) and 7 (on the trolley problem). Moreover, even within the narrower frame of the alternative actions available, the agent's permission to divert the trolley onto an innocent victim does not imply a permission to throw an innocent passenger onto the tracks to stop the trolley (as Thomson observes). Thus, even the immediate consequences (in terms of net lives saved) are not dispositive of permissibility.

[58] For one endorsement of such a threshold test, see Weinrib, *The Idea of Private Law*, 148–50. For a response, see Kenneth W. Simons, "Justification in Private Law" (book review of Weinrib), *Cornell Law Review* 81, no. 3 (March 1996): 711–12. Some proponents of threshold tests mean to endorse strict liability rather than negligence; my disagreement does not extend to them.

[59] But see Hurd, "Deontology of Negligence." For a critique, see Simons, "Deontology" (*supra* note 12), 290–95.

[60] Accordingly, many tort commentators view the "BPL" test and even the vaguer balancing test of the *Restatement (Second) of Torts* as necessarily utilitarian. See Richard Wright, "The Standards of Care in Negligence Law," in Owen, ed., *Philosophical Foundations*, 250 (listing sources).

be permitted to impose the risk, P × L must not only be greater than B, it must be *much* (or disproportionately) greater.[61] This could also be called a "thumb on the scale" test: in weighing the potential victim's interest in personal security against the potential injurer's interest in freedom of activity to impose risks, we should place a (heavy) thumb on the scale, giving special weight to the interest in personal security.[62]

These tests sound plausible and appealing, but, unless substantially recast, they provide a useless criterion. If we have identified the appropriate factors to balance, and if the method of balancing is also justifiable, then these tests say the following: One should not take a risk (as opposed to taking a precaution against the risk) simply because the advantages of taking the risk are greater than the disadvantages. Rather, one should take such a risk only if the advantages of doing so are *much* greater than the disadvantages (normally, only if the benefits to the injurer are much greater than the expected injuries to victims, discounted according to their probability).

This approach is either indeterminate or irrational. For unless one has a common metric or other justifiable method for measuring the competing interests or values, how does one know whether the interest in physical security and safety is "just" weightier than the interest in freedom of activity, as opposed to "much" weightier, so as to apply the "disproportion" or "much weightier" criterion? On the other hand, if one *does* have a common metric for measuring the competing interests, or if one does have some other justifiable method of balancing, why shouldn't the actor simply choose the alternative that furthers the "weightier" value, even if that value is only weightier by a peppercorn?

Let me be more specific. Is the interest in avoiding the risk of having one's arm broken "usually" greater than the interest in driving ten miles per hour faster, or "usually" greater than the interest in avoiding the expense of a more effective bumper?[63] These questions are meaningless unless we specify more clearly both the degree of risk of a broken arm, and the disadvantages of taking a precaution. Yet once we specify these factors, and adopt a justifiable method of balancing, shouldn't we indeed balance "at the margin"? That is, shouldn't we examine whether the

[61] Some have asserted that British courts employ the disproportion test, rather than a supposed simple cost-benefit American test, to gauge negligence in tort law. See Gilles, "Invisible Hand Formula," 1026 n. 8; and Gregory C. Keating, "Reasonableness and Rationality in Negligence Theory," *Stanford Law Review* 48, no. 2 (January 1996): 352–53. It is possible, however, that the British test merely shifts the burden of persuasion, and therefore does not create the problems noted in the text. See Gilles, "Invisible Hand Formula."

Ernest Weinrib appears to endorse some combination of the disproportion test and the "threshold of risk" test. See Simons, "Justification in Private Law," 702–4.

[62] See Keating, "Reasonableness," 354: "The *magnitude* of the harm that death, serious physical injury, and property damage threaten to persons' capacity to pursue their conceptions of the good is *usually* much greater than the *magnitude* of the harm threatened by increased precaution costs."

[63] See ibid.

advantages of any particular action (even a narrowly defined action) exceed the disadvantages?

I suspect that the worry about weighing "at the margin" is a legitimate concern about turning moral analysis into a bloodless form of calculation. What one should do should not depend on plugging numbers into a formula. And we should often be suspicious of methodologies that purport to balance along a "razor's edge,"[64] such that trivial factual differences in the weight of a given factor render an otherwise permissible action impermissible (or vice versa).

These concerns are well-founded if the most justifiable method of balancing requires a strong form of commensurability, that is, translation of all values into a single metric such as money or wealth. But weaker forms of commensurability are more plausible for most moral decisions, including decisions about risky alternatives.[65] For example, consider the question of whether a doctor should disclose to a patient all adverse risks of medical treatment of which the doctor is aware. There is a range of possible rules, from a rule of no disclosure (if the doctor believes that nondisclosure of a particular risk is in the best interest of the patient), to a rule of relatively full disclosure (of all risks that most patients would consider material), to a rule of disclosure tailored to the second-order preference of patients (that is, disclosure of whatever scope of risks the patient herself prefers to be disclosed).[66] These different rules embody different conceptions of the proper scope of patient autonomy and physician discretion in decision making about medical risks. Whether a given risk should be disclosed in a given case is much more likely to depend on these subtle value judgments than on the precise magnitude of the risk or on the precise financial or temporal burden to the doctor.

At the same time, however, even this more qualitative form of balancing will be sensitive to facts. Accordingly, close questions will sometimes arise about whether, for example, a particular risk is one that most patients would consider material. If we conclude that a doctor should disclose a 1 percent risk that hernia surgery will result in permanent numbness at the location of the surgery, but we find this a very close question, then

[64] See ibid., 353.

[65] Incommensurability between values A and B occurs in the strongest sense when adding to or subtracting from value A does not affect the choice between A and B. See Raz, *Morality of Freedom*, 322–26. (An example might be the question of whether Bach or Darwin was "more brilliant"; if either had been a little more or less brilliant than he in fact was, it would still be the case that neither was more brilliant than the other. Anderson, *Value in Ethics and Economics*, 55–56.)

Such incommensurability is rare. For example, if the choice between patient autonomy and burdens to a doctor of disclosing a risk were strongly incommensurable in this sense, then no increase or decrease in the burden to the doctor, or in the value of autonomy, would affect the choice. But that seems highly doubtful. In cases of risky activity, I believe, the competing values are, at most, incommensurable in this strong sense only within a limited range or "margin." See Raz, *Morality of Freedom*, 327–28.

[66] The British adopt a form of the first rule; most American jurisdictions adopt a form of the second; and some American jurisdictions permit the patient to waive disclosure, thus recognizing a version of the third.

the doctor might have no duty to disclose a 0.5 percent risk. In this sense, "marginal" decisions will still occur.

The "thumb on the scale" approach might also be designed to express special concern for one value in the balance, *relative* to some other, deficient way of valuing it. But this concern can be accommodated in a balancing test without suggesting the implausible conclusion that there will never be marginal cases. For example, one might conclude that the social value to be given to patient autonomy is greater than the value that most patients actually express in the marketplace (either because of marketplace distortions in capturing the private valuations of patients, or because recognizing patient autonomy is a collective social good, the value of which transcends the sum of individual valuations). Thus, even if patient surveys reveal that most patients only strongly care about risk information that has at least a 20 percent probability of changing their mind about treatment, the "thumb" might justify a rule that doctors disclose risk information with at least a 10 percent probability of changing a patient's mind.[67] Notice, however, that this use of a "thumb on the scale" is much more limited than the general use described earlier.

b. Balancing "freedom" against "security." Another proposal to systematize the nonconsequentialist balancing of values is the suggestion that the morality (and legality) of actions risking physical injury depends on a balance of "freedom" against "security"—the freedom of the injurer to engage in activity, compared to the interest of the victim in security against physical harm. A number of academics have supported some version of this balance,[68] and the theories of Immanuel Kant and John Rawls have been offered as justifications for such a balance.[69]

[67] For another example of a "relative valuation" approach, see Thomson, *Realm of Rights*, 197–99, where she argues that, to save A's life, a guardian can authorize a fairly serious nonfatal operation on A (such as cutting off A's leg), while a guardian of both X and Y cannot balance so close to the margin in authorizing an operation on X in order to save Y's life. More generally, insofar as extrapolating from an intrapersonal to an interpersonal case is feasible and defensible, we should at least apply a significant "premium" and should require a much greater total benefit in the interpersonal case than in the intrapersonal, in order to justify nonconsensual imposition of risks on others who do not receive any direct benefit from imposing those risks.

[68] See Keating, "Reasonableness," 319–27 (where he discusses balancing the injurer's freedom of action against the victim's interest in physical security); and Weinrib, *The Idea of Private Law*, 84–113 (where he discusses deriving tort principles from a paradigm of "doing and suffering"). Professor Richard Wright, in "Rights, Justice, and Tort Law," and "The Standards of Care in Negligence Law," in Owen, ed., *Philosophical Foundations*, describes corrective justice in tort law as expressing the Kantian notion that adversely affecting another's person or stock of resources is objectively inconsistent with the other's equal negative freedom. I agree with many of Wright's conclusions, including the inadequacy of the utilitarian account and its failure to explain the relevance of the actor's motive, the victim's consent, and other important factors. However, I am not persuaded that Wright's own reductionist framework can explain all the features of tort doctrine that he purports to explain. Rather, I believe that those features are more readily justified by a pluralistic analysis.

[69] See Keating, "Reasonableness"; Weinrib, *The Idea of Private Law*; and Wright, "Rights, Justice, and Tort Law."

This contrast is indeed helpful, but it is only part of the story. A number of values relevant to the permissibility of action do not fit easily into the categories of "freedom" and "security." One example is that just mentioned—patient autonomy in making decisions about medical treatment. If we reduce this value to "security," we ask only whether the patient's interest in decision making will in fact promote her physical safety. This approach ignores the possibility that some sacrifice in physical security is warranted in the interest of the patient's freedom to decide. If our only concern was the patient's physical well-being, we might support a significant degree of medical paternalism; but many would reject paternalism in favor of protecting a patient's ability to make decisions about her own bodily integrity.

c. Balancing: A pluralist approach. In the end, a pluralist form of balancing is the most attractive approach. What does a pluralist approach look like? Let me note some important features.

First, the actor will often be promoting social as well as personal purposes in his risky activity. Rather than broadly classifying all these purposes as aspects of an undifferentiated interest in "freedom of activity," we should attend to, and differentially value, the particular goods being sought. More virtuous motives, such as an altruistic concern for one's children, or for the welfare of a person in need of rescue, can be valued more highly than personal pleasure from the thrill of athletic effort. Serving a community's medical needs might have more weight than responding to people's consumption desires.[70]

Second, qualitative features of the risk and of the harm can be important. For example, is the risk fully understood by those in danger? Is it voluntarily incurred?[71] Are its benefits and burdens widely and equally shared? Does it conform to general expectations of the level of risk typical of the activity? With respect to harms, are they easily compensable (in money damages or otherwise)? Do they seriously interfere with the victim's long-term life plans?[72]

Third, rather than attending only to the advantages and disadvantages that the actor's risky activities foreseeably produce, we should attend also to the *reasons* why the actor was willing to take the risk, and to other features of the context in which the risk of harm was brought about. Insofar as the activity produced a risk of physical harm, was this risk or this harm intended, or merely an unfortunate side-effect? (Recall the "doc-

[70] See Anderson, *Value in Ethics and Economics*, 158–63.

[71] Whether explicit consent is a morally necessary precondition of imposing risk on others depends, in part, on the gravity of the risk. Below a certain threshold of risk, perhaps, no explicit consent is required. But the requirement of obtaining consent turns on the kind, as well as the level, of risk. Direct physical invasions of bodily integrity, such as through medical treatment, are far more likely to require explicit consent than are other risky acts with no greater, or even a lesser, risk of causing physical harm.

[72] See Keating, "Reasonableness," 344. However, I find problematic Keating's explicit reliance on a Rawlsian social-justice framework to justify private tort doctrine.

trine of double effect" examples, above.)[73] Did the actor at least make genuine efforts, even if inadequate, to reduce the risks? Was the risk part of a beneficial package that the actor offered to the victim? (Compare the risks to a product's consumer with the risks to a bystander injured by a product.)

One important aspect of how the risk is created is whether the actor has a special responsibility not to bring about wrongs, that is, an agent-relative duty not to act in certain ways, rather than an agent-neutral duty to minimize bad states of affairs. Consider here Robert Nozick's concern about a "utilitarianism of rights."[74] Nonconsequentialist norms sometimes forbid an actor from wronging a person, or violating his rights, even when the actor would thereby prevent others from committing more wrongs (or committing more rights-violations). For example, such norms forbid lying to prevent more lies, or convicting one innocent person in order to prevent more innocent people from being convicted.[75] The following example is more relevant to the topic of negligence: "A doctor may not neglect the health of her patient, a corporate executive whose demise will cause his firm to cease neglecting its workers' health."[76]

Fourth, the permissibility of a risky act sometimes depends on the social role of the actor. Professionals are properly held to special standards of skill and integrity; parents have special duties to their children; friends have special obligations to one another. Once again, it is doubtful that these special responsibilities can be reduced to concerns about "freedom" and "security."

Fifth, permissibility might also reflect basic distinctions between acting and omitting to act, or between harming and failing to confer a benefit. Thus, we might permit D to expose P to X amount of harm in order to avoid exposing P to Y amount of harm, so long as X is less than Y.[77] And yet we might not permit D to expose P to X amount of harm merely in order to provide a larger quantity Y of social *benefits*. (*A fortiori*, D may not expose P to X amount of harm merely to provide a larger quantity Y of *personal* benefit to D.)

[73] The actor's purpose or intention to bring about a harm has enormous significance. Even if such a purposeful actor reasonably believes that he is unlikely to succeed (and thus would otherwise be considered no more than negligent), he will usually be considered more culpable than an actor who believes he is likely to bring about the harm but does not intend it. See Simons, "Rethinking Mental States," 478–82.

[74] Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), 28. See also Raz, *Morality of Freedom*, 278.

[75] See, e.g., Thomas Nagel, "Personal Rights and Public Space," *Philosophy and Public Affairs* 24, no. 2 (Spring 1995): 89–90 (where Nagel relates this prohibition to the idea of an agent's "inviolability"). See also F. M. Kamm, "Non-consequentialism, the Person as an End-in-Itself, and the Significance of Status," *Philosophy and Public Affairs* 21, no. 4 (Fall 1992): 381–89.

[76] Anderson, *Value in Ethics and Economics*, 73.

[77] We might permit this even where Q rather than P is exposed to Y amount of harm, if P and Q are sufficiently similar in their vulnerability to an original threat of harm, as in a variation of the trolley problem.

In sum, when we evaluate the moral permissibility of risky action, we should draw qualitative as well as quantitative distinctions between the goods or values at stake. Just how those qualitative distinctions should be captured is a difficult question. Perhaps some values have an absolute "lexical priority" over others;[78] perhaps they "trump" other values;[79] perhaps they operate only as "side-constraints,"[80] or as claims with varying degrees of stringency,[81] or with "prima facie" weight.[82] But on any of these views, the consideration of competing values requires more subtle analysis than the simple maximization of aggregate value.

D. Hybrid (consequentialist and nonconsequentialist) views

Any plausible moral view will consider consequences. But there are several ways in which a moral view might include both consequentialist and nonconsequentialist concerns. First, as noted above, a modified utilitarian view might, for *nonconsequentialist* reasons, exclude certain preferences, pleasures, or values (such as the benefits a racist receives from indulging his preferences).

Second, nonconsequentialist concerns might serve as limits or side-constraints on achieving consequentialist goals. Thus, a general utilitarian approach to balancing the advantages and disadvantages of taking risks, and a utilitarian focus on preventing future acts of negligence, could be combined with moral side-constraints—for example, restricting blame to individuals who possess minimal general capacities to foresee risks and to avoid creating them.

Third, an apparently consequentialist balancing of the advantages and disadvantages of taking risks might reflect only a local rather than a global maximizing strategy. Suppose that the only reason for insisting that a moral agent maximize the expected costs and benefits of his action in choosing whether to take a risk is a retrospective, not a prospective, reason: creating unnecessary social costs is blameworthy. (As Richard Posner once put it, "we are indignant at a negligent injury because our moral natures are offended by economic waste, illustrated by an accident avoidable at a lower cost than the expected accident cost.")[83] But suppose the agent is not being blamed for failing to contribute to a broader, or global, maximizing strategy: in deciding whether to blame the agent, we

[78] See Rawls, *A Theory of Justice*, 42–45. Rawls describes a "lexical" ordering as one which "requires us to satisfy the first principle in the ordering before we can move on to the second, the second before we can consider the third, and so on."

[79] See Ronald Dworkin, *Taking Rights Seriously* (Cambridge: Harvard University Press, 1978), xi (where Dworkin argues that individual rights are political "trumps" held by individuals over collective goals [including utilitarianism]).

[80] See Nozick, *Anarchy, State, and Utopia*, ch. 3.

[81] See Thomson, *Realm of Rights*, chs. 6, 7.

[82] See W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930).

[83] Posner, *Tort Law*, 8.

pay no heed to the possible further beneficial consequences of blaming him (including deterrence of similar acts by him or by others). Then consequentialist considerations do not yet justify blaming the actor. For we have not followed those considerations to their reasoned limit. Instead, we have truncated the consequentialist strategy—either arbitrarily, or, more likely, because we have a nonconsequentialist reason for the limit. (For example, we might believe that actors should impartially consider the interests of victims and injurers, apart from the future benefits of employing such an attitude.)

Fourth, the *converse* type of hybrid is also defensible. Instead of adopting a (truncated) utilitarian criterion of negligence while ignoring potential deterrence, we might adopt a nonutilitarian criterion of negligence yet care about deterrence. Suppose, to take an extreme view, that the criterion of wrongdoing is this: it is impermissible to impose any risks on persons who do not fully consent in advance. This criterion is compatible with deterrence, if the main point of our blaming practice is to minimize future instances of nonconsensual risk-imposition.[84]

Which approach to combining consequentialist and nonconsequentialist concerns is most persuasive and normatively attractive? I will not attempt to resolve that question here, at the level of principle. But I do wish to emphasize that a number of different hybrids are normatively defensible, and that common-sense morality does reflect both consequentialist and nonconsequentialist concerns. Moreover, in a significant range of cases involving risks of physical injury, maximizing the value of good expected consequences is morally defensible, and indeed justifies the same decisions about risk that a nonmaximizing, nonconsequentialist balancing approach justifies. The question of how fast a person should drive, for example, normally does not present serious problems about the fair distribution of risk, about the types of human interest (such as sadistic preferences) that are morally valuable, about some interests having lexical priority over others, or about distinctive responsibilities attaching to special roles. Rather, the principal issues are the relative magnitude of the risks of injury, and the relative social benefits, of driving at differing speeds.

Now reconsider the (mythical) Ford Pinto case. The following hybrid approach is plausible. It is permissible for a product manufacturer to balance costs and benefits, so long as consumers are aware of the risks and are able to consent to them (through market decisions), and so long as the distributional effects are just. (Thus, we might consider the extent to which cost savings from not taking precautions will be passed on to vehicle users and potential victims.) A value must indeed be placed on

[84] For a similar analysis in the context of tort doctrine, see Gary Schwartz, "Mixed Theories of Tort Law: Affirming Both Deterrence and Corrective Justice," *Texas Law Review* 75, no. 7 (June 1977): 1828–33.

human life, for purposes of deciding what level of precaution the manufacturer should take;[85] but the "value" need not be a fixed figure, invariant in different contexts.[86] Also, the values to be balanced should include such qualitative ones as whether the risk of death by fire or explosion is especially dreaded. It is therefore highly unlikely that a mathematical calculation will be dispositive.

On the other hand, balancing risks of injury against social benefits is not always permissible. Consider the second example from the introduction, the example of Amy and Beatrice. Although Amy's creation of risks of physical injury as a byproduct of obtaining a possible health benefit for her own child is permissible, Beatrice's creation of similar risks as a byproduct of indirectly producing a diffuse future social good is impermissible. It matters how one brings about a benefit, not just how large the benefit is.[87]

Finally, the analysis to this point confirms the earlier suggestion that the differences between the negligent and "knowing" (or intentional) creation of harm are largely a question of degree. Although similar principles apply, the creation of lower-probability risks will, as a *factual* matter, more often be justifiable. For lower-probability risks more often require balancing of competing alternatives, more often trigger concern about consequences, and are less often subject to absolute or strong prima facie duties.

IV. Law

The relation of nonlegal standards of moral responsibility for negligence to legal negligence standards is a complex subject. Here I will only touch on a few points.

1. One function of legal standards is to reinforce moral standards. Criminal law and tort negligence standards do reflect and reinforce the moral standards discussed above. The moral duty not to act negligently is an important source of a variety of legal rules, including tort-law injunctive remedies for nuisance, tort-law damages for the harms brought about by negligence, criminal-law sanctions for risky behavior and for the harms brought about by such risky behavior, and governmental regulation of risk. Indeed, when legal responses to risk seem directly to *contradict* conventional moral standards, popular opposition is sometimes remarkably intense. An example is the difficulty that many states have encountered replacing tort fault liability for certain automobile accidents with a no-

[85] On a deontological view, this value need not be equal to the amount provided as compensation after the fact.

[86] See Pildes and Sunstein, "Reinventing the Regulatory State," 43–86.

[87] For an illuminating discussion of this feature of deontological reasoning, see Leo Katz, *Ill-Gotten Gains* (Chicago: University of Chicago Press, 1997).

fault standard, notwithstanding the enormous financial costs of the former system.[88]

2. Legal institutions obviously have enormous advantages over private moral standards in addressing a host of problems, including collective action, inadequate information, the difficulty of coordinating one's actions with those of others, and insufficient incentives to comply with moral norms. Coercive legal institutions permit risks to be reduced in widely varying ways (from police protection, to traffic rules, to facilitating economic markets), and to a degree that would otherwise be unachievable.

3. The moral duty of reasonable care itself depends in part on legal rules. For example, if we focus only on isolated individual responsibility for risk-creation, we ignore the point that the risks that I may fairly impose on others depend on what others will do, or are entitled to do. Yet what others will or may do is often significantly determined by legal rules—including property and contract law, and laws that coordinate activities (such as traffic rules).

4. The distinctive features of legal institutions help explain some distinctive features of legal conceptions of negligence. Consider the legal construct of the "reasonable person." Moral theory has no need for such a construct; we could, in principle, directly articulate the relevant principles for permissible and impermissible action, and for culpable and nonculpable behavior, without such an ideal. But the constructive legal approach can be very different from the direct approach.

At first glance, the distinction between (a) demanding a good or morally sufficient justification for imposing a particular risk, and (b) asking whether a reasonable person would impose that risk, might seem trivial. It is not. Depending on how one constructs a "reasonable person," this evaluation could be quite different from the direct evaluation of justifications. For example, the analysis of the justifiable use of defensive force could directly specify the relevant factors (one may not use more force than the amount of force one is defending against; one may defend only where necessary to prevent a future harm). Or the analysis could instead employ a "reasonable person" criterion, specifying merely that acceptable force should be judged by the standard of when, and to what extent, force would be used by a "reasonably prudent person, cautious about his own and others' safety, and not overwhelmed by emotion," or something of the like. The latter standard is likely to produce a less determinate set of results than the former standard, and could produce either a narrower or a broader self-defense privilege.

[88] Despite much academic support for automobile no-fault insurance as a replacement for tort negligence liability, and despite evidence that no-fault schemes save considerably on the enforcement costs of negligence liability, many people are offended by the idea that someone can negligently damage their car or their body without paying directly for that harm. (Of course, the opposition of trial lawyers is also an important part of this story.)

A "reasonable person" test often serves an important institutional function: one institution lays down some general constraints or permissions, while another specifies the details via a "reasonable person" test. This institutional division of labor permits a trier of fact to develop normative standards case by case, instead of having a legislature define them *ex ante*. (Thus, a legislature might define criminal or tort negligence simply as "lack of reasonable care," leaving a jury or other fact-finder with the considerable task of spelling out the meaning of that standard.) Absent this institutional differentiation, a "reasonable person" test would have little use.[89]

Sometimes, to be sure, the "reasonable person" construct is merely a convenient shorthand for moral considerations that could be spelled out if one had the time or inclination to do so. Thus, it is not surprising that reasonableness standards are often used in describing *epistemic* duties, that is, duties to form reasonably accurate beliefs. What a "reasonable person" would believe or foresee could be specifically articulated in terms of types and strength of evidence, relevant purposes, and the like; but the inquiry is sufficiently complex, and there is sufficient consensus about what people "should" believe, that the shorthand is acceptable. The same might be true when the question is whether an actor conducted an activity with adequate skill. Thus, it is convenient to refer to the skill that a reasonable driver would exercise in braking, or making a turn; but, in principle, we might be able to spell out the factors relevant to a "reasonable" exercise of skill.[90]

5. Legal institutions are capable of doing much more about risks of harm than simply reinforcing private norms against unjustifiable risk-taking. Government can (and should) address risks that are no individual person's responsibility. Consider, for example, risks of harm created by nature, not by persons. Government might have a duty to forecast bad weather and other natural disasters, and to conduct medical or safety research.

[89] This division of labor also allows the formal legal standard to remain constant through time, while flexible in its application. (I thank Hugh Baxter for this point.)

Whether a virtue-theory approach can give content to a "reasonable person" standard is an interesting question, but one that I do not have the space to explore here. On virtue theory generally, see Daniel Statman, ed., *Virtue Ethics* (Washington, DC: Georgetown University Press, 1997).

[90] Insofar as a "reasonable person" test is meant to establish a standard that people can fairly be blamed for not satisfying, the test should be at least partially subjective—i.e., it should assess the individual capacities of the agent. See Perry, "Risk, Harm, and Responsibility," 344. Legal standards are sometimes more "objective" than this, both for practical reasons (avoiding problems of proof and of fraud) and also, perhaps, in order to express a (strict liability) principle of fairness (e.g., the principle that others in the community are entitled to a relatively high standard of conduct, even if this requires blaming or holding liable some who cannot reasonably be expected to meet the standard). Thus, courts typically ignore intellectual deficiencies, and also religious convictions that prompt actors to impose higher than usual risks on others (or on themselves).

On the other hand, the collective and coercive features of law clearly make it an inappropriate instrument for enforcing all private moral norms. The range of human actions that are negligent and deserve modest moral blame far exceeds the range of actions that properly deserve legal redress through tort or regulatory sanction; and the range of actions that properly deserve criminal condemnation is much narrower still. The law should not be concerned with trivial risks of harm, or with forms of neglect and errors of judgment that are common and understandable. An automobile owner's failure to abide by the manufacturer's maintenance schedule is not on a par with drunk driving. Citizens should have the liberty to be modestly negligent, especially insofar as the law seeks to prevent or enjoin the act of negligence, as opposed to compensating for its ill effects.

6. If law simply enforced the moral norms that were customary in a community, or that were supported by most respected philosophers (even if such consensus were possible!), law would obviously risk losing its legitimacy. The content of legal norms should have a democratic pedigree. Private moral norms obviously differ in this respect: they are not "enacted" through a democratic process. To this extent, the model of law reinforcing ordinary moral judgments is problematic. For example, if a common-law court in a tort case saw its function as enforcing private moral norms (and if legislatures could not overrule that choice), tort doctrine would have only an attenuated grounding in democratic choice.

The legitimacy problem is profound. Some would invoke this difficulty to support a utilitarian approach, relying on utilitarianism's apparent neutrality between different preferences, its indifference to the sources and relative value of preferences, and its associated anti-perfectionism. But utilitarianism itself is clearly a controversial moral perspective. Although some versions of utilitarianism might succeed in not taking sides at the ground level of individual conceptions of the good, utilitarianism most certainly takes sides on the "higher-level" questions of how to resolve conflicts among such conceptions and how to define the social, rather than individual, good.

To resolve these difficulties, one needs a defensible political theory that explains which norms can justifiably be legally enforced in light of the variety of moral views people hold in a pluralist society. Developing such a theory is not easy, and I certainly will not attempt it here.[91]

V. Tort Law, Fault, and Corrective Justice

Tort law is a useful context in which to test some of the above analysis. The tort law construct of a "reasonable person" illustrates some distinc-

[91] The difficulty that John Rawls has encountered in convincing skeptics of the value of his own attempted solution underscores the seriousness of this problem. See John Rawls, *Political Liberalism* (New York: Columbia University Press, 1996); and Rawls, "The Idea of Public Reason Revisited," *University of Chicago Law Review* 64, no. 3 (Summer 1997): 765–807.

tive features of law, while the *ex ante* concept of negligence helps illuminate the debate between "corrective justice" and "fault" interpretations of tort doctrine.

1. In Anglo-American tort law, the typical jury instruction explaining the meaning of negligence refers only to the conduct exercised by a largely undefined "reasonable person" or "reasonably prudent person in the circumstances."[92] Such an instruction invites juries to identify and articulate the moral norms of the community.[93] At the same time, courts have developed distinct doctrines for such problems as causation and the scope of duty (for example, duties of landowners, special relationships, no general duty to rescue). Often these more specific doctrines reinforce community moral judgments, but sometimes they reflect the distinctive role and limits of legal institutions. For example, ordinary morality surely imposes at least some minimal duty to rescue a stranger in need, but tort law refuses to impose a general duty to rescue. An important justification for that refusal is a concern to protect personal liberty from state regulation, even the rather incidental form of state regulation that tort liability imposes.

Moreover, although the general "reasonably prudent person" test might seem to be without substantive content, courts do give greater definition to the concept of negligence in appellate rulings, and in some particular doctrinal areas, especially product liability. Their definitions vary, but they do support some of the nonutilitarian moral conceptions of negligence explained above.[94]

2. An important question for the interpretation and justification of negligence doctrine is whether we treat *fault* or *corrective justice* as a more fundamental concept. By "fault," I mean unjustified conduct, including negligence and unjustified intentional harms. By "corrective justice," I mean the Aristotelian idea that an agent has a duty to rectify a harm that he has caused, when it would be unjust to leave the harm unredressed.[95]

[92] See Gilles, "Invisible Hand Formula," 1017.

[93] See Michael Wells, "Scientific Policymaking and the Torts Revolution," *Georgia Law Review* 26, no. 3 (Spring 1992): 731; and Catherine Pierce Wells, "Tort Law as Corrective Justice: A Pragmatic Justification for Jury Adjudication," *Michigan Law Review* 88, no. 8 (August 1990): 2348-2413.

[94] Sometimes, jury instructions explaining negligence reflect a duty of impartiality, of considering interests of others as you would consider your own. See *Restatement (Second) of Torts*, section 283 cmt. e (1965); Keating, "Reasonableness," 337-38; and Gilles, "Invisible Hand Formula," 1038. For further discussion of the moral content of negligence doctrine, see Simons, "Deontology," 277-85.

[95] The criticism that follows might not apply to much broader interpretations of corrective justice, such as Margaret Radin's understanding of corrective justice as any principles governing our response to the wrongful or unjust unsettling of entitlements. See Margaret Jane Radin, "Compensation and Commensurability," *Duke Law Journal* 43, no. 1 (October 1993): 60. On the other hand, insofar as Radin's view refers only to *ex post* "correction" of a wrong, it might not fully encompass the *ex ante* perspective of negligence.

Most recent academic writings that defend a nonutilitarian account of tort law treat corrective justice as more basic than fault. I believe that this is a mistake. (It is also a mistake that I have committed in the past.)

As an interpretation of existing tort doctrine, the corrective-justice view understates the significance of fault, and overstates the significance of strict liability principles. Anglo-American tort law does not begin with the presumption that a loss should be shifted to the injurer. Rather, it normally does not shift a loss unless the loss was caused by faulty conduct. The categories in which strict liability is recognized are limited;[96] even product liability, which according to much judicial rhetoric is strict, is predominantly fault-based.[97]

To be sure, nominally fault-based tort doctrines might contain elements of strict liability that are better explained by corrective justice. The objective test of negligence, for example, sometimes sets a standard that the defendant cannot fairly be expected to meet. This severity might reflect utilitarian concerns about fraud and difficulty of proof. But it might also reflect a particular corrective-justice rationale: that one who is victimized by another's excusable failure to live up to the normal standard of care in the community is entitled to compensation (even if she cannot fairly complain that the other was at fault and should have acted differently).

Also, as a practical matter, insofar as hindsight bias is pervasive, triers of fact in tort cases will often greatly overestimate the *ex ante* probability of harm, once they know (as they always do) that the harm has actually occurred. Accordingly, the actual application of a purported *ex ante* negligence test probably results in a substantially stricter form of liability than negligence doctrine can justify. (A conscientious judge therefore should advise a jury to avoid such hindsight bias.)

Still, in the end, fault principles explain Anglo-American tort doctrine better than corrective-justice principles do. Moreover, fault principles also offer a better *justification* than corrective justice for most of tort law. A principal reason is the *ex ante* perspective of fault. An actor's primary duty is not to create unreasonable, unjustifiable risks of injury. The duty to compensate or to provide another form of redress, in case the actor actually causes harm, is distinctly secondary. If it were possible to *enjoin* negligent acts, before they could result in harm, we would often do so. It is simply a fortuity of the natural world that most individual acts of negligence are too isolated, random, and unpredictable to be prevented *ex ante*; but the logic of fault suggests that they should often be prevented,

[96] These include: abnormally dangerous activities (such as the use of explosives), wild animals, product liability to some extent, and vicarious liability. (The latter imposes strict liability for the tort of another; but that other tort is usually fault-based.)

[97] Liability for defective design and defective warnings is largely fault-based. Liability for manufacturing defects is strict, however.

if that were feasible.[98] By contrast, nonfault principles of strict liability do indeed suppose that harm is actually done; they impose a *primary* duty to compensate for such harm.[99]

Concededly, the corrective-justice interpretation seems to be supported by the following feature of tort law: normally the same damages (full compensation) are available regardless of the degree of the actor's culpability. Even a strictly liable defendant, who might not be culpable at all, pays full compensation. (Compare this to criminal law, in which the penalty varies with the culpability or dangerousness of the actor, not just with the harm done.) This might suggest that negligence reflects a corrective-justice orientation toward redressing the harm done to the victim, not an orientation toward imposing duties on those who have acted with fault. Still, the occasional availability of punitive damages shows a clear concern with fault, transcending the concern to redress a loss. And in any event, the prevailing practice of normally awarding only damages for harm done reflects pragmatics, not a principled preference for damages over injunctions or other forms of prevention when (as is rarely the case) the latter responses are feasible.[100]

A corrective-justice theorist might further reply by pointing out an especially attractive feature of his approach: the fact that it distinguishes corrective justice from distributive justice. Corrective justice is not concerned with whether the victim was entitled to the property that the injurer harmed, or to the bodily integrity that the injurer intentionally or accidentally invaded. Rather, it presupposes an initial set of entitlements, and asks when the injurer should compensate the victim or provide some

[98] Indeed, in nuisance law, if a person's use of his property is unreasonable, and if certain other criteria are satisfied, injunctive relief is presumed to be the proper remedy.

The qualification "often" leaves room for a liberty constraint on state power when the moral fault of the defendant is modest. If an actor drives drunk, it is permissible for the state to prevent or enjoin his conduct. If a driver fails to follow the automobile manufacturer's maintenance schedule, enjoining him to do so would be an excessive use of state power. In either case, however, requiring compensation for harms caused by the actor's neglect might be a legitimate use of state power. (See the discussion in the prior section.)

A corrective-justice theorist might reply that compensatory liability in such a case exemplifies strict liability rather than fault. But (as I have argued elsewhere) a genuine fault approach could justify imposing an *ex ante* tort "fine" to approximate the expected costs of the neglect (in either of the cases mentioned in the previous paragraph). Kenneth W. Simons, "Corrective Justice and Liability for Risk-Creation: A Comment," *U.C.L.A. Law Review* 38, no. 1 (October 1990): 113–42. If this is correct, then *ex post* compensation can also be viewed as fault-based.

[99] I do not address here "strict liability" principles of *conditional* fault, e.g., a requirement that a business provide insurance or some other *ex ante* assurance that it will be able to pay for harms that it might cause. See Simons, "Jules Coleman and Corrective Justice," 880.

[100] Indeed, it would be justifiable to employ a different measure of damages in strict liability cases than in negligence cases—for example, routinely adding a "kicker" to negligence damages (beyond those damages that would leave the victim indifferent between damages and harm) to reflect the special wrong of creating unreasonable risks to others. For similar suggestions, see Keating, "Reasonableness," 349 n. 125, and sources cited there.

other appropriate remedy for unjustly disturbing those entitlements. What entitlements people should have is the subject of a distinct mode of justice, distributive justice.

But a fault analysis can, and indeed should, draw the same distinction, and should not attempt to resolve the distinct issue of the initial justice of entitlements. At the same time, a fault analysis has this advantage over corrective justice: it can more easily condemn behavior that risks, but has not yet caused, harm to entitlements.

VI. Conclusion

Several important themes emerge from this essay. First, "negligence" is both an important concept and an ambiguous one. Of its many meanings, I have concentrated upon the sense of creating an unjustifiable, low-probability risk of future harm. Second, I have attempted to dispel the prevalent view that only a maximizing, utilitarian approach can render intelligible certain features of negligence analysis—its focus on the *marginal* advantages and disadvantages of the actor's taking a specific precaution, its consideration and balancing of the short-term effects of different actions, and its sensitivity to a multiplicity of factors. Perhaps certain absolutist deontological perspectives are inconsistent with these features; but other deontological perspectives (not to mention other nonutilitarian and partially nonconsequentialist perspectives) can easily accommodate them. Third, I have tried to show how these moral perspectives are reflected in legal rules, and particularly in tort doctrine. Careful examination of the concept of negligence helps resolve an important debate about the nature of tort law, supporting the view that fault, rather than corrective justice, is the better interpretation and justification of Anglo-American tort doctrine.

In some ways, the ambitions of this essay have been modest. I have not attempted to conclusively justify a definitive set of negligence principles that should govern us as a matter of private morality or as a matter of legal enforcement. But I hope that I have succeeded in showing that conventional moral understandings, and plausible moral principles, contradict the reigning utilitarian account of negligence.

Law, Boston University