

STATISTICAL KNOWLEDGE DECONSTRUCTED

Boston University School of Law Working Paper No. 10-26
(September 7, 2010)

Kenneth W. Simons

Revision of March 1, 2011

This paper can be downloaded without charge at:

<http://www.bu.edu/law/faculty/scholarship/workingpapers/2010.html>

Statistical Knowledge Deconstructed

by

Kenneth W. Simons*

Draft: March 1, 2011

Abstract:

The law frequently distinguishes between individualized knowledge (awareness that one's act will harm a particular victim, e.g., driving through an intersection while aware that one's automobile is likely to injure a pedestrian) and statistical knowledge (awareness that one's activity or multiple acts will, to a high statistical likelihood, harm one or more potential victims, e.g., proceeding with a large construction project that one confidently predicts will result in worker injuries). Under tort and criminal law doctrine, acting with individualized knowledge is ordinarily much more difficult to justify, and, if unjustified, much more culpable, than acting with statistical knowledge. Yet the distinction is very difficult to explain and defend.

In this article, the first systematic analysis of this pervasive but underappreciated problem, I offer a qualified defense of the distinction. Acting with statistical knowledge is ordinarily less culpable than acting with individualized knowledge, and often is not culpable at all. Expanding the spatial or temporal scope of an activity or repeating a series of acts sometimes causes the actor to acquire statistical knowledge, but such an increase in scale ordinarily does not increase the level of culpability properly attributable to the actor. Two invariant culpability principles, "Invariant culpability when acts are aggregated" and "Invariant culpability when risk-exposures are aggregated," formalize this idea.

Why is acting with individualized knowledge especially culpable? Part of the answer is the special stringency principle (SSP), a deontological principle that treats an actor as highly culpable, and treats his acts as especially difficult to justify, when he knowingly imposes a highly concentrated risk of serious harm on a victim. Under SSP, speeding to the hospital to save five passengers, knowing that this will likely require killing a pedestrian in one's path, is much harder to justify than speeding to the hospital to save one passenger, knowing that this creates a 20% chance of killing a pedestrian in one's path.

The analysis has a number of significant implications but is also subject to important qualifications:

- Notwithstanding the invariant culpability principles, if a faulty or unjustified actor repeats his acts or expands his activity, that repetition sometimes reveals a new type of culpability: the defiance of moral and legal norms. Accordingly, a retributivist can indeed support a punishment premium for recidivists.
- In rare cases, when the actor possesses merely statistical knowledge but his conduct is extremely unjustifiable, the actor's culpability is comparable to that of an actor with individualized knowledge.
- The higher culpability of acting with individualized knowledge is not explained by a supposed higher duty owed to "identifiable victims," except insofar as that duty is a crude version of SSP.
- When an actor conducts a cost-benefit analysis of a planned activity and thereby acquires statistical knowledge that the activity will cause serious harm, his decision to proceed with the activity despite that knowledge is not, by itself, evidence of his culpability.
- A legal system can be legitimate even though legal actors within the system know that it will, as a statistical matter, punish the innocent.

* Professor of Law & The Honorable Frank R. Kenison Distinguished Scholar in Law, Boston University School of Law. © Kenneth W. Simons 2010; all rights reserved. For helpful comments, I thank Matt Adler, Larry Alexander, Eric Blumenson, Bob Bone, Mike Cahill, Alon Harel, Kim Ferzan, Leo Katz, Adam Kolber, Gerry Leonard, Mike Meurer, Nancy Moore; participants at Boston University School of Law, Brooklyn Law School and [University of Texas School of Law, March 2011=] faculty workshops; and participants at the Harvard Law School Private Law Workshop. I am also grateful to Joseph Cooper, Andrew Keutmann, Greg Racki and Renee Williams for excellent research assistance and editorial advice.

Outline

- I. Introduction
- II. The legal relevance of mental states including cognitive mental states
 - A. Mental states in general
 - B. Cognitive mental states
- III. Statistical knowledge deconstructed
 - A. Specifying the problem
 - B. A partial solution: the principle of ICAA (Invariant culpability when acts are aggregated)
 - C. Does ICAA apply to faulty actors?
 - D. The special stringency principle (SSP) for highly concentrated risks
 - E. Is the real question whether the victim is identifiable?
 - F. Why are we tempted to equate statistical with individualized knowledge?
 - G. How should we distinguish statistical from individualized knowledge?
 - 1. The principle of ICREA (Invariant culpability when risk-exposures are aggregated)
 - 2. Statistical knowledge of a risk that is extremely unjustifiable
- IV. Applications
 - A. Knowledge in contexts other than knowingly causing serious harm
 - B. If a cost-benefit analysis supplies knowledge that an activity will cause harm, is it culpable to proceed?
 - C. Can retributivists tolerate the predicted mistaken punishment of the innocent?
 - D. Practical and doctrinal issues
- V. Conclusion
- VI. Appendix

I. Introduction

The law typically treats an actor who knows that he will likely cause serious harm, and who does cause that harm, as a very culpable wrongdoer, meriting a severe sanction. Consider a simple illustration:

Clara (a speeding driver who acts with individualized knowledge)

Clara is running late for a plane to a critically important business meeting. Speeding along a narrow and deserted mountain road, she suddenly sees a drunk young man, barely moving, lying on the road in the path of the car. Rather than stop for him, she slows down but continues, striking the man as her car passes over him. She realizes that she is very likely to injure him, which she does.¹

Clara is likely to be liable in tort for a knowing battery, with ensuing punitive as well as compensatory damages, and for criminal assault and battery as well. And if she realizes that she is likely to kill the victim, and does kill him, she will very likely be guilty of murder.

But the Anglo-American legal system does not treat actors so severely in cases where the knowledge is (what I will call) “statistical” rather than, as in Clara’s case, “individualized.”² Such comparative leniency is ordinarily the correct approach, as we will see. Yet it is remarkably difficult to explain and justify.

Consider two examples of actors who possess statistical knowledge that they will cause harm:

Agatha (a careful builder who acts with statistical knowledge)

Agatha is building an enormous tunnel.³ Based on past experience with projects of this size and scope, she can confidently predict that even with the highest level of care currently attainable, a dozen of the thousands of workers who will be employed over the ten year period of the project will

¹ The illustration is loosely based on a famous hypothetical from Sandy Kadish’s criminal law casebook. See S. Kadish, S. Schulhofer, & C. Steiker, *Criminal Law and Its Processes* 842 (8th ed. 2007).

² Other discussions of the problem also use the term “statistical” knowledge. See Sanford Kadish, *Respect for Life and Regard for Rights in the Criminal Law*, 64 *Calif. L. Rev.* 871, 891-894 (1976); = Armour, Henderson. Thomas Schelling appears to have originated the terms “statistical life” and “statistical death,” to be contrasted with losses of identified, individual lives. Thomas Schelling, *The Life You Save May Be Your Own*, in Thomas Schelling, *Choice and Consequence* 115 (1984). See Gregory C. Keating, *Pricelessness and Life: An Essay for Guido Calabresi*, 64 *Md. L. Rev.* 159, 160 n. 4 (2006).

³ For simplicity, assume that she is the owner of the construction company and makes all relevant decisions about taking precautions. This obviates the difficult question of assessing the mental state of an organization or corporate body. (Worker’s compensation laws will affect the ability of the workers to obtain tort recovery from Agatha; for purposes of this discussion, assume that they are a category of workers not barred from tort recovery by such laws.)

suffer serious injuries, and, as a statistical matter, there is a 90% chance of at least one death. Agatha does employ a high level of care. A dozen workers do suffer a serious injury, as predicted. One dies.

Knowingly injuring a person counts as battery under both criminal law and tort law. Is Agatha guilty of a dozen counts of criminal law battery? Is she liable to the dozen victims for tortious battery? For punitive damages? Knowingly causing a death usually counts as murder. Is she guilty of murder of the worker who died? Under current law, the answer to all of these questions is no.

Now consider this variation:

Bertha (a less careful builder who acts with statistical knowledge)

Bertha is building an enormous tunnel, of the same size and scope as Agatha's. But Bertha takes less extensive precautions than are customary, based on advice that these lesser precautions are almost as effective as the usual ones, and are less costly and time-consuming, permitting the tunnel to be completed much earlier, with resulting social benefits. The projected injury rate is approximately 30% higher than would be expected for Agatha's project.

Bertha might be found negligent, or perhaps reckless, in deciding to use lesser precautions. But it is quite unlikely that she will be found guilty of a knowing or intentional battery if the lesser precautions cause some additional workers to suffer injuries. And it is quite unlikely that she will be found guilty of murder if her lesser care resulted in the death of one more worker than she would statistically expect if she had used due care.

Are these disparate legal results correct? If so, how can they be explained? Or conversely, if in the second scenario Bertha should be deemed a "knowing" injurer or killer, why shouldn't Agatha be so deemed in the first scenario (even though, by hypothesis, she used all reasonably feasible care)? Moreover, recall Clara, the speeding businesswoman who will be considered a knowing wrongdoer, under both tort and criminal law doctrine. Is the disparate legal treatment of Bertha and Clara justifiable? If so, why?

Finally, consider an example from prominent criminal law theorist Sanford Kadish⁴ that vividly highlights the contrast:

⁴ Kadish, *supra* note 2=, at 892. I have slightly modified his example. In the original, five to fifteen lives were almost certain to be lost despite the exercise of due care, and the cost of not lowering the fitting in place was an indefinite but not necessarily permanent abandonment of the tunnel.

Louis Kaplan and Steven Shavell offer a very similar example in their famous article. Louis Kaplow & Steven Shavell, *Fairness versus Welfare*, 114 *Harv. L. Rev.* 961, 1276 n. 757 (2001) (last para. of footnote). However, Kaplow and Shavell infer from the example, not that there is a distinction between individualized and statistical knowledge, but that on a consequentialist view, we should draw no distinction; apparently they would not object to the decision of Eleanor to sacrifice the worker.

Eleanor (a builder who acts with individualized knowledge)

Eleanor is building an enormous underwater tunnel. Based on past experience with projects of this size and scope, she can confidently predict that even with the highest level of care currently attainable, three to five lives are almost certain to be lost by the end of the project. Early in the construction process, she discovers that a worker is trapped in a section of the partially laid tunnel. Unfortunately, a fitting must be lowered into place within an hour, or else the whole tunnel will have to be permanently abandoned due to changing river conditions; and one hour is insufficient time to rescue the worker. She knows that laying down the fitting will surely crush the worker to death. She proceeds to lower the fitting, and the worker dies.

Kadish observes: “I expect that it would nonetheless be a form of criminal homicide to lower the fitting. Even if it were justified under a lesser-evils formula, which is doubtful, the decision would be a soul-searching one. Yet attaining the very same social good—the construction of the tunnel—readily justified its construction despite the predicted loss of multiple lives.”⁵ Our cast of characters reinforce Kadish’s observation: the law will treat Eleanor very differently from Agatha or Bertha. Again, the question is: why?

In this article, I defend a qualified legal (and moral) distinction between individualized knowledge (exemplified by Clara and Eleanor) and statistical knowledge (exemplified by Agatha and Bertha). But clarifying, justifying, and sensibly qualifying the distinction are no easy matter, as we will see.

On the one hand, there are strong *prima facie* reasons to reject the idea that one who acts with statistical knowledge that he will cause harm should be considered extremely culpable (and just as culpable as one who acts with individualized knowledge). Consider the following reasons:

- **“Scale should not matter”**: If I drive carefully on a single occasion, imposing low risks on others, I do not deserve legal sanction if I happen to cause harm. Why, then, should I suddenly become liable merely because I choose to drive in the same manner every day? If I hire one worker, after carefully checking his documentation, I do not deserve legal sanction if it turns out that he is an undocumented immigrant. Why, then, should I suddenly become liable for knowingly hiring an undocumented immigrant⁶ if I employ 100,000 workers and I know that my careful screening will nevertheless result in my hiring at least one undocumented worker?

Below, I formalize the argument that multiplying (a) similar acts or (b) similar occurrences of risk-exposure, either over time or space, does not by itself change the culpability of the underlying act. The principle corresponding to (a) is “ICAA,” shorthand for “Invariant Culpability when Acts are Aggregated”; corresponding to (b) is “ICREA,” or “Invariant Culpability when Risk-Exposures are Aggregated.”⁷

- **“Careful risk/benefit analysis should not inculcate”**: Many believe that sound and justifiable regulatory and tort law policy requires companies to

⁵ *Id.* at 892.

⁶ “It is unlawful for a person or other entity ... to hire ... for employment in the United States an alien knowing the alien is an unauthorized alien.” 8 U.S.C. §1324a (a)(1)(A).

⁷ See TAN = *infra*.

conduct careful risk-benefit cost-benefit analyses of whether to adopt precautions that would reduce the health and safety risks of their products or activities. It would be inconsistent with this desideratum then to impose sanctions on companies merely because they conducted such analyses and carefully examined the probable consequences of their actions. But if statistical knowledge should be treated as harshly as individualized knowledge, it follows that a company that chooses to engage in far-flung activities or chooses to manufacture or distribute a large number of products should be treated as highly culpable simply because (1) it conducted a careful cost-benefit analysis and (2) the scale of its activities is vast. That is a most unwelcome conclusion. And it is a conclusion explicitly rejected by provisions of the recently adopted Restatement Third of Torts.⁸

- **“The modern hierarchy of mental states requires a distinction between statistical and individualized knowledge”**: On the conventional modern view of the hierarchy of mental states, knowingly causing harm is roughly comparable in culpability to purposely causing harm; recklessly causing harm is less culpable; and negligently causing harm is even less culpable. But if statistical knowledge is really just as culpable as individual knowledge, this framework crumbles, because many cases of statistical knowledge seem indistinguishable from cases of recklessness, of negligence, or even of faultless conduct.
- **“Knowingly harming another is distinctively wrongful only when the actor’s knowledge is individualized”**: Consider two distinct versions of this crucial argument:
 - (1) **“We owe a higher duty to identifiable victims”**: When the actor has statistical knowledge, the victim is typically unidentified, while if he has individualized knowledge, the victim is identifiable. And we arguably have a more stringent moral and legal obligation to save an identifiable victim than to take advanced precautions against harm to unidentified possible future victims. Governments, and coal companies, spend much more (per life saved) to save trapped coal miners than they spend to prevent coal mining accidents in the first instance. This distinction reflects political reality but arguably is also justifiable, in light of our more stringent duty to save those we see in immediate peril than to prevent future harm to the large and indeterminate group of victims that we calculate will be at risk if we take inadequate safety precautions.

⁸ Restatement of the Law Third, Torts: Liability for Physical or Emotional Harm, §1, comment e, page 9:

[I]n many situations a defendant's knowledge of substantially certain harms is entirely consistent with the absence of any liability in tort. For example, an owner of land, arranging for the construction of a high-rise building, can confidently predict that some number of workers will be seriously injured in the course of the construction project; the company that runs a railroad can be sure that railroad operations will over time result in a significant number of serious personal injuries; the manufacturer of knives can easily predict that a certain number of persons using its knives will inadvertently cut themselves. Despite their knowledge, these actors do not intentionally cause the injuries that result. Moreover, despite their knowledge, none of the companies—absent further facts—can even be found guilty of negligence; nor does the knowledge possessed by the knife manufacturer subject it to liability under products-liability doctrines.

For further discussion, see TAN = infra.

(2) “We owe a special duty not to impose highly concentrated risks of harm on a single person”: On this view, it is wrongful to concentrate the costs of an activity on one person merely to provide small benefits to many other persons. Nonconsequentialists object that utilitarians are too willing to justify acts whose aggregate benefits outweigh their costs, even when the costs are concentrated on a few persons or even a single person, while the benefits are widely diffused. Consider Greg Keating’s version of philosopher T.M. Scanlon’s famous World Cup example:

Suppose that a piece of transmitting equipment has toppled and crushed a television technician helping to broadcast [the World Cup live] to a billion viewers worldwide, and that the only way to save the technician’s life is to interrupt the broadcast for thirty minutes, effectively thwarting the transmission of the show⁹

The objection is powerful and persuasive: clearly it would be wrong to knowingly permit the technician’s death simply in order to provide a small benefit to each viewer; and this is so even if the number of viewers is vast. But then the question arises: how can we permit other risky activities, such as high-speed (but careful) driving, when they are justified only by the small and diffuse benefits they provide to those engaged in the activity? A partial answer is that when an actor drives carefully, even at high speed, on a single occasion, he does not have individualized knowledge that someone on the road will suffer death or bodily injury.

And yet, on the other hand, there are also strong reasons to question the view that those who act with merely statistical awareness that they will cause harm are less culpable than those who act with individualized knowledge. Consider these arguments:

- **“Statistical victims are still victims”:** The distinction between statistical knowledge and individualized knowledge appears to rest on a fallacious claim that if a victim is not identifiable in advance, if he is a mere statistic, then his death somehow counts for less. But as Professor Lisa Heinzerling has argued, “[s]imply calling what we are valuing ‘risk’ . . . rather than ‘life,’ does not change the fact that real lives, not statistical lives, hang in the balance. There are no statistical people. There are only real people.”¹⁰ And it is arguable immoral, and even irrational, to care more about “identifiable” victims than about other victims. More bluntly: to permit the predictable killing of “statistical” victims is to permit “statistical murder.”¹¹
- **“Juries often reject the distinction”:** Juries often view in a very harsh light a company’s deliberate decision not to adopt a safety precaution when the

⁹ Gregory Keating, *Pressing Precaution Beyond the Point of Cost-Justification*, 56 *Vand. L. Rev.* 653, 666 (2003). I have changed Keating’s original example from a broadcast of a popular television show to a broadcast of the World Cup.

¹⁰ Lisa Heinzerling, *Knowing Killing and Environmental Law*, 14 *N.Y.U. L. Rev.* 521, 531 (2006). See also Lisa Heinzerling, *The Rights of Statistical People*, 24 *Harv. Env. L. Rev.* 189, = (2000); Guido Calabresi, *Ideals, Beliefs, Attitudes, and the Law* 6 (1983) (“I have never run into a ‘statistical’ life anywhere, in my car or anyplace else.”).

¹¹ See James Hammitt & Nicolas Treich, *Statistical vs. identified lives in benefit-cost analysis*, 35 *J Risk Uncertainty* 45, 45 (2007), citing Joanne Linnerooth, *Murdering Statistical Lives...?*, in Michael Jones-Lee (ed.) *The Value of Life and Safety* (1982), p. 231. See also Allan Gibbard, *Risk and Value*, in *Values at Risk* (Douglas MacLean, ed. 1986), pp. 94-112, at 97 (wondering why this is not “murder for hire”).

company knows the likely harmful consequences of taking less care. This condemnation suggests that juries at least sometimes consider statistical knowledge morally equivalent to individualized knowledge. The lawsuits involving the decision of Ford not to strengthen the Pinto's fuel tank, and of General Motors to place the 1979 Malibu's fuel tank close to the rear bumper, arguably fit this pattern. In each case, the jury was presented with evidence of the company's statistical knowledge that their design would cause more injuries and deaths than an alternative design, and in each case, the jury imposed considerable punitive damages, beyond ordinary compensatory damages.¹²

- **“Ordinary people often find cost-benefit analysis abhorrent”**: Experimental research has explored the behavior of mock juries confronted with the decision whether to sanction a company in the following three scenarios:
 - (a) The company conducted an unreasonable cost-benefit analysis of whether to adopt a precaution;
 - (b) The company conducted a reasonable cost-benefit analysis;or
 - (c) The company acted unreasonably but failed to conduct any cost-benefit analysis.

Jurors are more likely to sanction with punitive damages companies in category (b) than those in category (c), even though companies in (b) are not negligent while companies in (c) are negligent.¹³ These studies arguably support a conclusion that causing harm while possessing statistical knowledge of the harmful effects of one's decision not to adopt a precaution (even when the decision itself is reasonable) is viewed as quite culpable—and more culpable than negligent ignorance of those effects. Perhaps such jurors believe that acting with statistical knowledge (acquired through the deliberate use of cost/benefit analysis) is about as culpable as acting while possessing individualized knowledge.

- **“Acting with merely statistical knowledge is sometimes still highly culpable”**: Some actors who arguably fit the “statistical” knowledge category rather than the “individualized” knowledge category are nevertheless highly culpable. Consider:

¹² In the 1978 Ford Pinto case, the initial punitive damage award was \$125 million (an extraordinarily large sum at that time), later reduced to \$3.5 million. In the 1999 GM Malibu case, the initial punitive damage award was \$4.9 billion, later reduced to \$1.1 billion; the case was subsequently settled for an undisclosed amount. Peter Hong, Judge Cuts Award Against GM to \$1.2 Billion, L.A. Times, Aug. 27, 1999, at B1. See W. Kip Viscusi, Corporate Risk Analysis: A Reckless Act?, 52 Stan. L. Rev. 547, 547, 569 (2000); Peter Y. Hong, Judge Cuts Award Against GM to \$1.2 Billion, L.A. Times, Aug. 27, 1999; Staff Writers, GM to Settle Case Over Gas Tank Explosion, L.A. Times, July 25, 2003. Accessed at: <http://articles.latimes.com/2003/jul/25/local/me-briefs25.1>. See generally Gary Schwartz, The Myth of the Ford Pinto, 42 Rutg. L.Rev. 1013 (1991).

¹³ See Viscusi, id. at 588. Although Viscusi's study is illuminating, one possibly legitimate reason for the participants' reactions is their arguably justifiable distaste for actors who engage in an especially narrow form of cost-benefit analysis. Another, less legitimate reason is hindsight bias: when an observer knows that the actor's risky conduct caused a serious harm, the observer is likely to overstate the risk that the actor should have recognized ex ante. See Robert Macoun, The Costs and Benefits of Letting Juries Punish Corporations: Comment on Viscusi, 52 Stan. L. Rev. 1821 (2000).

Time bomb

A terrorist plants a bomb in a central train station. He knows that the bomb is constructed in such a way that it will explode at an indeterminate time in the near future, indiscriminately wounding or killing many individuals who happen to pass by.

Should we not conclude that the terrorist knowingly kills? Is he not as culpable as the other characters we have already encountered who possessed individual knowledge, Clara and Eleanor? And yet it is indeterminate in advance who will be killed, and when they will die.

This last example suggests that perhaps other actors who possess statistical knowledge, who know that their activities will result in death at an indeterminate time, to indeterminate victims, and at indeterminate locations, should similarly be treated very harshly. If a manufacturer knows that a small number of his products contain defects that will cause harm, should he not be treated like the terrorist? Isn't his product a "ticking time bomb"? And similarly for the tunnel builders Agatha and Bertha; isn't the decision to embark on a huge construction project similar to triggering a bomb or distributing a poison that will have destructive effects at some point, to someone?

The conundrum of statistical knowledge is of both theoretical and practical interest, and has been a focal point of controversy in a very wide range of legal contexts. This article is the first systematic examination of the issue. Solving the conundrum, or at least reducing the confusion it engenders, illuminates a range of different questions, including:

- (a) The basic rationale for treating individualized knowledge as an especially culpable state of mind;
- (b) Whether recidivists deserve a punishment premium;
- (c) Whether one owes a more stringent duty not to harm "identifiable" victims;
- (d) Why many people intuitively (but often wrongly) condemn corporate decision-makers for proceeding in the face of a cost-benefit analysis which predicts that their decision not to take a safety precaution will cause death or serious harm; and
- (e) Why a legal system can be legitimate even though it predictably punishes the innocent.

The article is organized as follows. Section II addresses the legal relevance of mental states in general, and cognitive mental states in particular. Section IIIA explains the scope of the problem. Section IIIB introduces the principle of "invariant culpability when acts are aggregated" (ICAA) as a novel analytical device, while IIIC examines whether the principle applies to faulty as well as faultless actors. Section IIID elucidates SSP, the special stringency principle. Many nonconsequentialist endorse an especially stringent duty not to impose concentrated risks on another, a duty that helps explain the presumptive unjustifiability of acting with individualized knowledge. Section IIIE asks whether the real issue in distinguishing individualized from statistical knowledge is whether the victim is identifiable; it answers in the negative. Section IIIF investigates why we are tempted to make the false equation of statistical knowledge and individualized knowledge.

Section IIIG considers some possible criteria for distinguishing statistical from individualized knowledge, and at the same time (1) delineates the principle of ICREA (Invariant culpability when risk-exposures, rather than acts, are aggregated) and (2) points out that in exceptional circumstances—namely, when the risk is extremely unjustifiable—statistical and individual knowledge are indeed normatively indistinguishable; but these circumstances actually “prove” the rule, i.e. they actually underscore the legitimate rationales behind ICAA and ICREA. Section IV applies the analysis to several issues, including: contexts where knowledge pertains to a legally relevant circumstance rather than a causal result; objections to cost-benefit analysis; and a serious objection to retributivism. The latter objection is that retributivists cannot justify any real-world system of punishment, because retributivists find unjustifiable the knowing infliction of suffering on those who do not deserve it, and yet any realistically imaginable legal system will predictably result in the mistaken punishment of the innocent. Section IVD attends to some practical and doctrinal issues. It is followed by a conclusion that responds to the questions I have raised in this introduction.¹⁴ An Appendix includes ICAA, ICREA, and SSP, and offers some paradigm individual knowledge scenarios, together with some intriguing variations.

II. The legal relevance of mental states including cognitive mental states

A. Mental states in general

As background to the analysis to follow, it is important to outline briefly some of the principal reason why mental states matter to legal liability.

¹⁴ I should briefly note some limitations and assumptions of this article’s analysis. First, in speaking of an actor’s beliefs in a particular probability that a harm will occur or that a circumstance exists, I am employing an epistemic and Bayesian conception of risk, not a frequentist conception. This conception is, I believe, the one most appropriate to the legal and moral issues explored here. See Matthew Adler, *Against ‘Individual Risk,’* 153 U. Pa. L. Rev. 1121, 1226-1227 (2005). However, a frequentist analysis of risk might be relevant to certain issues, such as loss of a chance damages in tort law. See Matthew Adler, *Risk, Death, and Harm,* 87 Minn. L. Rev. 1293, 1436-1444 (2003).

Second, I do not examine here the distinction between individual and population risk. See Adler, *Against ‘Individual Risk,’* *supra* at 1126. I doubt that that distinction has much relevance in typical criminal law and tort law scenarios. See Kenneth W. Simons, *Tort Negligence, Cost-Benefit Analysis, and Tradeoffs: A Closer Look at the Controversy,* 41 Loyola L. Rev. 1171, 1218-1221 (2008). Or, more precisely, I doubt the significance of the distinction when the risks in question are relatively small. However, when the individual risk is especially high, the difference is indeed significant for a nonconsequentialist. Suppose D1 creates a 90% risk of killing each of 10 people, while D2 creates a 0.9% risk of killing each of 1,000. D1 and D2 create the same population risk, for in either case, the expected harm is 9 deaths. However, D1 alone creates an extremely high individual risk to the group he endangers, triggering what I call the Special Stringency Principle, discussed *infra* at . This justifies a higher legal sanction for D1 than for D2 (assuming that the motives, benefits, and other relevant features of their acts are otherwise the same).

Third, the statistical knowledge conundrum that I consider pertains to the actor’s culpability for “knowingly” causing harm or for “knowing” that a legally relevant circumstance (such as the illegal status of his employee) exists. Thus, I am not addressing the evidentiary problems raised by the “blue bus” hypothetical and similar quandaries, where we are unsure whether one actor or another caused the harm, and where we have only probabilistic evidence (e.g., “60% of the buses in this town are owned by Blue Bus Co.”) to resolve the question. A valuable recent discussion of these evidentiary issues is Mike Redmayne, *Exploring the Proof Paradoxes,* 14 Legal Theory 281 (2008).

An actor's state of mind is frequently a crucial determinant of his legal responsibility. It can determine his legal responsibility *vel non*, as well as the extent or grade of liability (and the availability of defenses). A criminal defendant's mental state can be the difference between guilt or innocence, or between mandatory life imprisonment and a much shorter prison term. A tort defendant's mental state determines whether she has committed an intentional tort or instead a tort of negligence, and that distinction entails further consequences—for example, whether victim fault may reduce recovery, or whether resulting harm will be deemed sufficiently “proximate.” Moreover, a tortfeasor is liable for punitive as well as compensatory damages only if she is “reckless” (in some sense), not if she is merely negligent. Finally, the mental state of a government actor is often relevant to liability, and sometimes decisive.¹⁵

These examples underscore the extraordinarily wide range of contexts in which mental states figure in legal doctrine. Mental state distinctions play a more central role, and are more extensively elaborated, in some legal contexts (such as criminal law) than in others. What is a bit surprising, though, is this: insofar as legal responsibility does depend at least in part on mental states, a single conventional hierarchy is pervasively employed. That hierarchy identifies four basic mental states reflecting decreasing legal responsibility: purpose (or intention), knowledge, recklessness, and negligence.¹⁶ The hierarchy has an attractive simplicity, and legal decision-makers often employ a full or partial version of it across doctrinal categories. The Model Penal Code is the exemplar, though tort law and other fields of law also use most or all of these categories.

A closer look reveals that the four culpability terms in the conventional hierarchy do not simply identify different degrees of culpability along a single spectrum; they also combine different qualitative dimensions of culpability, dimensions of belief, desire, justifiability, and deviation from a normative standard.¹⁷ Thus, purpose or intention requires a “conscious object”¹⁸ to bring about a result (for example, in murder, the death of a victim), and thus typically requires both a belief that the result is possible and a desire (in some sense) that the result

¹⁵ Thus, a government actor is subject to liability for intentional discrimination on the basis of race or sex, but not for knowingly imposing disproportionate harm upon a racial minority group or women (under the Equal Protection clause); for deliberately restricting speech because of disapproval of the speaker's ideas, but not for imposing the same restriction on the basis of a legitimate, and not censorial, policy (under the First Amendment); and for deliberate indifference to the medical needs of prisoners, but not for mere negligence (under the Eighth Amendment).

¹⁶ The hierarchy also recognizes a fifth no-fault or strict liability category, identifying cases in which the actor has none of the four basic types of mental state but still is subject to legal liability.

It is fair to ask whether negligence is a genuine “mental state.” In many contexts, it is used as a criterion of substandard behavior, without any reference (or any direct reference) to a state of mind. See Larry Alexander, *The Philosophy of Criminal Law*, in *The Oxford Handbook of Jurisprudence and Philosophy of Law* 815, 829, ch. 20 (J. Coleman & S. Shapiro eds. 200=); Kenneth W. Simons, *Rethinking Mental States*, 72 *B.U. L. Rev.* 463, = (1992). But for simplicity, in this paper I will describe it as a mental state, rather than as a “culpability term,” a phrase that is the more accurate but less familiar.

¹⁷ See Simons, *supra* note 16=.

¹⁸ Model Penal Code §2.02(2)(a)(i).

occur.¹⁹ Knowledge (insofar as it is simply a mental state²⁰) requires a belief that a result is quite likely to occur.²¹ Recklessness, in a common modern formulation, requires both (1) a belief that there is a significant or substantial risk that the harmful result will occur, and (2) that the taking of the risk is unjustifiable.²² And negligence is often defined as culpable inadvertence to risk; it requires a deviation (or a gross deviation) from the standard of care that a reasonable person would observe, either in appreciating a risk or in acting to avoid or minimize the risk.

The application of these four culpability terms to circumstances, such as a victim's nonconsent, rather than results, such as the harm to a victim, is analyzed in a similar fashion²³: we can ask whether an actor who intentionally had sexual intercourse with a victim also intended (or desired or hoped) that she not consent, whether he knew she was not consenting, whether he was reckless as to that fact, or whether he was negligent.

Consider four homicide examples, in each of which the victim dies.

Don causes death purposely, knowingly, recklessly, or negligently

If Don deliberately crashes his car into Victor with the intention of killing him, he causes the death purposely. If he drives his car through a crowd while escaping from the police, recognizing that he is very likely to kill one of the pedestrians but not desiring or intending to do so, he causes the death knowingly. If he drives very near a pedestrian, recognizing that he might hit and kill him, but believing that he most likely will not do so, he causes the death recklessly. And if Don is engrossed in reading Professor Simons' most recent work on mental states while absent-mindedly driving in stop-and-go traffic, and therefore does not see a pedestrian walk in front of his car, he causes the pedestrian's death negligently.

¹⁹ Here I elide some important complexities. For further discussion, see Michael Bratman, *Intentions, Plans, and Practical Reason* (1987); F.M. Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (Oxford 2006).

The MPC definition of purpose as "conscious object" does not specifically require that the actor believe that the means he chooses have some possible chance of achieving the result, but this is probably implicit in the quoted phrase. (If you shoot a gun at an astronaut on the moon, desiring his death, but believing it impossible to achieve this goal, you probably don't have the "conscious object" of killing him, though you do hope for that result.)

²⁰ A legal requirement that D know that X actually combines a mental state requirement, that D believe that X, with an objective requirement, that X is true. In order for an actor to knowingly cause a harm, for example, the actor must believe that the harm will occur, and the harm must indeed occur. The latter requirement is not a mental state. For convenience, however, I will follow convention and characterize the mental state as "knowing" rather than "believing." (Among other things, there is no similar linguistic equivalent of the adverb "knowingly." If I believe that I will cause someone's death yet fail to do so, I cannot have acted "knowingly." Rather, I acted "believingly." The latter term, alas, is not (yet?) in any standard English dictionary.)

²¹ MPC §2.02(2)(b), 2.02(7); Restatement (Second) of Torts §8A; =Cite to torts restatement 3rd.

²² MPC §2.02(2)(c).

²³ But not identically. See TAN = below.

In this essay, which is part of a larger project of critiquing the reigning hierarchy of mental states, I focus on, and question, one aspect of the hierarchy—the law’s treatment of cognitive mental states, and particularly the mental state that tort and criminal law typically call “knowledge.” Keep in mind that this term encompasses only a belief that a result is highly likely to occur (or that a circumstance is highly likely to exist). “Knowledge,” then, is a term of art: it departs from standard usage insofar as we often do colloquially speak of “knowledge” that one’s conduct creates a small risk of harm (“She knew that her dangerous driving might kill someone”), or of “knowledge” that there is a small risk that a fact (such as the victim’s lack of consent) exists (“He knew that she might not have consented to their sexual relations.”). These latter examples, where the actor is aware of a small risk rather than a high likelihood of harm, are instances of cognitive “recklessness,” not knowledge, under modern doctrine.

B. Cognitive mental states

The culpability of knowledge is much more complex and interesting than is generally acknowledged.²⁴ The argument in this article will concentrate on one particular context in which a culpability of knowledge is especially problematic—when a person knows that his repeated conduct or extensive activity will very likely, in a statistical sense, cause harm.

Before we examine the statistical knowledge conundrum, it is useful to consider more carefully how cognitive mental states like knowledge, belief, and recklessness are relevant to legal responsibility.

A legal requirement of knowledge always includes a requirement that the actor have some type and degree of subjective belief. A requirement of recklessness usually also includes a requirement of subjective belief.²⁵ To know that a result will occur is (inter alia) to believe that it will surely, or very likely, occur; to be reckless is to believe that the result has a significant or substantial chance of occurring. In principle, in judging the degree of an actor’s responsibility, the law could employ belief requirements spanning the entire range of probabilities, from a belief that the probability is 100% to the belief that it is just above zero.²⁶ As a practical matter,

²⁴ Two examples:

(1) Does “knowingly causing harm” have the same meaning and significance when, instead of acting, I have failed to act (when I had a duty to do so)? Is a lifeguard who simply panics and thus fails to rescue a drowning swimmer a knowing killer, i.e., a murderer?

(2) Suppose I know a legally relevant fact but, when I later act and cause harm, I have forgotten the fact. Should I be treated as still knowing it?

²⁵ However, sometimes recklessness criteria do not require subjective awareness; instead, they require either gross negligence or some form of culpable indifference. To avoid unnecessary complications, I do not discuss these variations.

Belief is also sometimes relevant to a negligence standard, especially in criminal law. The Model Penal Code essentially defines negligence as unreasonable failure to be aware of a substantial risk, i.e., failure to believe what a reasonable person would have believed. MPC §2.02(2)(d). But in other contexts (and most of the time in tort law), belief is not relevant in this way to negligence: an actor’s behavior can be negligent because it falls below an idealized standard of care (such as a standard of proficiency or skill), not because, or not merely because, it is based on negligent beliefs.

²⁶ I am discussing inculpatory belief requirements; there would be no point to a “requirement” that the actor believe that the risk of harm from his action is zero. But as a matter of exculpation, such a belief is of course highly relevant. An actor who inflicts deadly force in self-defense (even in circumstances where this is disproportionate to the threat) but

however, legal standards pick out at most only two or three points on the continuum. Thus, with respect to circumstance elements, the Model Penal Code distinguishes awareness of a substantial risk, awareness of a high probability, and simple belief that the circumstance exists.²⁷ Some other criminal codes only identify one point, and tort law standards often identify only one, or perhaps two.²⁸

We should also remember that “recklessness,” as defined in the influential Model Penal Code and in many state criminal codes, is a complex concept, combining two requirements. The actor must both (a) believe that the risk of the relevant legal fact is substantial and (b) take a risk that society considers unjustifiable. In most of this essay, I will be using the term in its simpler, purely cognitive sense, i.e., as encompassing requirement (a) but not requirement (b). Thus, unless otherwise noted, the reader should assume that “reckless” is an instance of cognitive recklessness, not complex recklessness. (This distinction is crucial, we shall later see.)

What is the significance of determining that an actor has or has not satisfied a mens rea requirement of knowledge? That requirement serves three crucial functions:

- (1) Grading offenses or legal wrongs, or varying sanctions, according to relative seriousness;
- (2) Establishing a threshold below which the conduct is not legally impermissible at all; and
- (3) Demanding an especially strong justification before the conduct will be deemed legally permissible (especially strong, that is, relative to the justification needed in the case of a lesser mental state such as recklessness or negligence).

A mental state of cognitive recklessness or knowledge can serve all three functions (and so can other mental states, such as intention).

First consider the grading function. In criminal law, a more culpable state of mind can lead to greater punishment. Cognitive mental states are especially apt means for achieving this grading function, because beliefs are scalar, i.e., differentiable according to the actor’s degree of certainty or confidence. Thus, acting with knowledge is often treated as more culpable than acting with the “lower” belief of cognitive recklessness. Not surprisingly, knowingly causing death is

who reasonably believed that there was zero or very little risk that his violent response would be that harmful will obtain a full defense based on his reasonable mistake.

²⁷ =cites. It treats the latter two degrees of belief essentially the same, as two different formulations of “knowledge.” Specifically, awareness of a high probability suffices for knowledge “unless [the actor] actually believes that [the fact] does not exist.” §2.02(7). For some difficulties with this formulation, see Kenneth W. Simons, *Should the Model Penal Code’s Mens Rea Provisions Be Amended?*, 1 Ohio State J. Crim. L. 179, = (2003). With respect to results, the MPC defines knowledge as belief that the result is “practically certain.” §2.02(2)(b)(ii).

²⁸ Under the Restatement (Second), intent is defined to include both purpose and knowledge. Negligence does not require actual awareness. Recklessness is an infrequently used mental state in tort law, and it is not clear whether it requires awareness of risk. See Simons, *Rethinking Mental States*, supra note 15=, at 483. And it is similarly unclear whether the reckless disregard often required for punitive damages, although sometimes defined differently than the recklessness occasionally used as a basic mental state for liability, requires awareness of risk. See *id.* at 484.

usually murder, while recklessly causing death is usually manslaughter.²⁹ With intentions, by contrast, for the most part either one has an intention or one does not; further legal gradations of degree are rarely invoked.

In tort, the significance of greater culpability for liability is less direct, since compensatory damages are measured the same way for both intentional torts and torts of negligence (or of strict liability). Still, characterizing a tort as “intentional” (which encompasses both purposeful and knowing infringements) aids the plaintiff in a number of ways—in sometimes permitting recovery for emotional harm even without physical harm; in precluding consideration of victim fault, in many jurisdictions; in relaxing requirements of proximate cause; and in more readily permitting the award of punitive damages.

Why do we so readily assume that a person who knowingly causes harm is more blameworthy or responsible than a person who recklessly causes harm? Because, in the usual case, if a person chooses to create a risk while believing that a harmful consequence is likely, he is, *ceteris paribus*,³⁰ more responsible or culpable than a person who creates a risk believing only that a harmful consequence is possible; for the first actor reveals, by his action, willingness to tolerate a greater harm or evil in pursuit of his ends. When we condemn an actor or hold her responsible for knowingly or recklessly causing harm, we are not condemning her simply for having a “guilty” or culpable state of mind as to results (or even, in the case of recklessness, for failing to fully recognize the seriousness of risks). Rather, or more precisely, we hold her responsible for her willingness to act notwithstanding her belief that her actions will or might cause harm.³¹ What is morally and legally significant is her decision to act in light of her beliefs.

²⁹ This is not a pure example of scalar grading, however. Jurisdictions (such as those following the Model Penal Code) that define manslaughter as “recklessly causing death” mean “recklessness” in the complex sense, as (a) cognitively reckless awareness of a risk (b) that is unjustified. But the second element, lack of justification, is not part of the definition of knowing murder; rather, necessity and self-defense remain as distinct defenses to murder. So comparing knowing murder with (complexly) reckless manslaughter is comparing apples with oranges. If reckless manslaughter were understood in the cognitive sense (element (a) only), then the comparison would be apples to apples.

³⁰ The *ceteris paribus* condition is important. Some cases of knowingly causing a harm are less culpable than some cases of recklessly (in the cognitive sense) causing a harm. See Larry Alexander and Kimberly Kessler Ferzan, with Stephen Morse, Crime and Culpability: A Theory of Criminal Law 33 (2009):

[S]omeone who imposes a very high risk of harm on another – a risk just short of practical certainty – for a very frivolous reason, such as a thrill, is surely more culpable than one who imposes a practically certain harm for a quite weighty, but ultimately insufficient, reason. For example, if Albert knowingly inflicts severe bodily injury on another to prevent him from accidentally destroying Albert’s entire life’s work, Albert may be culpable; but he is surely less culpable than if he imposes a slightly lower risk of the same harm on another just to satisfy his urge to drive like a madman. In short, instances of what we now call extreme indifference to human life may not only be equal to but also more culpable than some cases of knowledge.

Moreover, if a person desires or intends to cause harm, it might be irrelevant whether he believes that the harm is very likely or merely believes that the harm has a substantial chance of occurring. See Itzhak Kugler, Direct and Oblique Intention in the Criminal Law: An Inquiry into Degrees of Blameworthiness 91-99 (2002). Under the Model Penal Code, “purpose” requires that the actor have the conscious object to bring about a result but does not require the belief that the result will occur with any specified probability. 202(2)(a)(i).

³¹ In the case of omissions, we hold her responsible for a failure to act when she has a duty to do so, and when she is aware that her action will or might prevent harm.

Second, consider the threshold function of a mental state. Knowledge requirements often serve this function. This is true not only of many crimes (for example, knowledge is often required for drug and other possession offenses), but also of many other types of wrongs. Thus, such “intentional” torts as trespass and battery can be committed only by purposely or knowingly invading the victim’s property interest or interest in physical integrity, not by invading such interests inadvertently.³² And, as a constitutional matter, tort liability for defamation of a public figure or public official requires proof that the defendant “recklessly disregarded” the truth; that phrase has been interpreted as requiring a state of mind close to knowledge.³³

Third, consider the justificatory function of mental states: actions accompanied by more “serious” or culpable mental states require more compelling justifications, and often only permit justifications in narrower categories, relative to actions accompanied by less “serious” mental states. In criminal law, a purposeful or knowing killing is murder and can only be justified by a narrow and particularly compelling set of reasons, such as self-defense or necessity.³⁴ But one who is or should be aware that his conduct poses a substantial risk of death, and is thus merely cognitively reckless or negligent, is only guilty in the first instance if the risk of death he poses is unjustifiable. Moreover, justification is much easier to demonstrate here than in the case of purposeful or knowing killing. A surgeon can be justified in posing even a very high risk of death, if the expected benefit of the operation is to save the patient’s life; but an automobile driver can be justified in posing a risk of death for a relatively trivial reason (e.g. he takes a drive in order to feel the wind rushing through his hair), so long as the risk of death he poses is relatively small.

These explanations of three crucial functions of mental states have thus far been stated somewhat abstractly. I will put some flesh on the bones of this analysis below, when I introduce several new characters—Larry, Maxwell, and Ned—and explore whether these functions operate differently in the contexts of individualized and statistical knowledge.

III. Statistical knowledge deconstructed

A. Specifying the problem

Recall the two earlier examples involving tunnel builders Agatha and Bertha. Careful Agatha could predict that even with a high (and non-negligent) level of precaution, injuries and even death will occur to a few of the thousands of workers over the lengthy duration of the project. Less careful Bertha could predict that with a lower (and probably negligent) level of precaution, injuries and death are even more likely. And recall Clara:

³² Nor is it the case that negligent invasions of those interests are necessarily actionable under the tort of negligence. The latter tort requires that the victim suffer physical harm; the torts of trespass and battery do not. Thus, purpose or knowledge indeed does serve a threshold function in trespass and battery. See Kenneth W. Simons, *A Restatement (Third) of Intentional Torts?*, 48 *Ariz. L. Rev.* 1061, 1080-1090 (2006).

³³ See discussion at notes = *infra*.

³⁴ Indeed, many states do not permit a necessity defense to murder. Joshua Dressler, *Understanding Criminal Law* = (5th ed. 2009); Heinzerling, *Knowing Killing*, *supra* note 10=, at 528. And when they do recognize necessity, they often limit the necessity to “emergencies.” Dressler, at =.

Clara (a speeding driver who acts with individualized knowledge)

Clara is running late for a plane to a critically important business meeting. Speeding along a narrow and deserted mountain road, she suddenly sees a drunk young man, barely moving, lying on the road in the path of the car. Rather than stop for him, she slows down but continues, striking the man as her car passes over him. She realizes that she is very likely to injure him, which she does.

The law will characterize Clara as “knowingly” causing that injury, but it will not characterize either Agatha or Bertha as “knowing” that she will cause death or serious injury to others.³⁵ And yet in one sense Agatha and Bertha do satisfy the usual criteria for knowingly causing harm: they subjectively believe that harm is almost certain to occur as a result of their activities. How can their mental states be distinguished from Clara’s? Or from the mental state of Eleanor, the builder who knowingly crushes the victim to death? Or must we swallow the (very) bitter pill and concede that all of these cases must be treated the same?

A very similar statistical knowledge problem arises when the object of the required knowledge is a circumstance that makes the actor’s conduct illegal, rather than (as in the examples thus far) a result of the actor’s conduct. This version of the problem is much less discussed, but it seems in principle the same. Recall the earlier example of a large corporation charged with the crime of “knowingly” hiring undocumented immigrants as employees. And suppose the corporation is so large, with hundreds of thousands of employees, that it can confidently predict that it will hire some illegal employees, even if it uses extraordinary efforts to comply with the law. Interpreting this criminal prohibition to embrace statistical knowledge would result in conviction for acts that are, by hypothesis, not even negligent. This implication makes it highly doubtful that courts would interpret the prohibition so broadly.³⁶ But why, exactly, would that broad interpretation be improper?

³⁵ There is little explicit case law on the topic. However, some tort cases have held that (what I call) statistical knowledge does not satisfy the intent requirement for battery. See *Shaw v. Brown & Williamson Tobacco Corp.*, 973 F.Supp. 539, 548 (D.Md.,1997):

Brown & Williamson did not know with a substantial degree of certainty that second-hand smoke would touch any particular non-smoker. While it may have had knowledge that second-hand smoke would reach some non-smokers, the Court finds that such generalized knowledge is insufficient to satisfy the intent requirement for battery.

See also *Madden v. D.C. Transit Systems, Inc.*, 307 A.2d 756 (D.C. 1973) (rejecting battery claim for lack of requisite intent where plaintiff claimed that defendant knew that its buses were regularly spewing fumes and oily substances).

Some academics have supported tort liability for statistically known harms. See generally Darren S. Rimer, *Secondhand Smoke Damages: Extending a Cause of Action for Battery Against a Tobacco Manufacturer*, 2 Sw. U. L. Rev. 1237, = (1995) (arguing that statistical knowledge should suffice for battery liability); Paul LeBel, *Intent and Recklessness as Bases of Products Liability: One Step Back, Two Steps Forward*, 32 Ala. L. Rev. 31, 67 (1980) (arguing that statistical knowledge should suffice for products liability if defendant is permitted a risk-utility defense).

³⁶ My research has not uncovered cases containing the claim that statistical knowledge suffices for legal liability when the fact the defendant “knows” is a circumstance rather than a result.

The statistical knowledge conundrum is also part of the contentious debate over the legal relevance, and even the legitimacy, of cost/benefit analyses conducted by corporate actors. Should we fault such actors simply for performing a deliberate, conscious cost-benefit calculus before engaging in risky activity? To do so is very troubling; after all, on at least some accounts of negligence and recklessness, actors have a moral and legal duty to conduct just such an analysis.³⁷ To be sure, some believe that a utilitarian comparison of costs and benefits is the wrong way to distinguish faulty from faultless behavior.³⁸ But even if that is so—indeed, even if conducting a purely utilitarian analysis is sufficiently wrong to deserve blame or legal sanction—it is an enormous further step to equate the culpability of an actor who makes such an analysis with the very serious culpability characteristically displayed by actors who cause harm with the mental states of purpose or individualized knowledge. And yet, once again, it is not clear why, in applying a doctrinal knowledge requirement, a corporate actor that can reliably predict that its activities, services, or products will cause harm should not be treated as a knowing actor, for such a prediction seems to satisfy the knowledge criterion.

A good way to begin answering the statistical knowledge conundrum is to revisit the three critical functions (noted earlier) that are served by characterizing a mental state as “knowing”: (1) grading or sanctioning the actor as especially responsible or culpable; (2) establishing a threshold, below which the action is not legally impermissible; and (3) demanding an especially strong justification before his conduct will be deemed legally permissible.³⁹ Consider how the law will treat statistical knowledge differently from individualized knowledge with respect to these three functions.

First, in criminal law, causing death with individualized knowledge is often punished more harshly than recklessly causing death: the first is typically murder, while the second is typically reckless or involuntary manslaughter. But a case of unreasonably risky conduct accompanied by statistical knowledge, such as the Bertha scenario, will at worst be treated as reckless manslaughter, not as murder.⁴⁰ Moreover, in tort law, punitive damages will often be available in individualized knowledge cases, but they will not normally be available in statistical knowledge

³⁷ See Viscusi, *Corporate Risk Analysis*, supra note 12, at 588 (reporting results of mock juror analysis showing that “undertaking any type of risk analysis was harmful to the corporation’s prospects both with respect to the probability of punitive damages and, more importantly, with respect to the magnitude of the award”); Henderson & Twerski =.

³⁸ Deontological views of negligence include George Fletcher, *Fairness and Utility in Tort Theory*, 85 *Harv. L. Rev.* 537 (1972); Gregory Keating, *Reasonableness and Rationality in Negligence Theory*, 48 *Stan. L. Rev.* 311 (1996); Kenneth W. Simons, *Negligence*, 16 *Soc. Phil. & Pol.* 52 (1999); Simons, *The Hand Formula in the Draft Restatement (Third) of Torts: Encompassing Fairness As Well as Efficiency Values*, 54 *Vand. L. Rev.* 901 (2001); Richard Wright, *Hand, Posner, and the Myth of the “Hand Formula,”* in *Symposium, Negligence in the Law*, 4 *Theoretical Inquiries in Law* 145 (2003); John Oberdiek, *Towards a Right Against Risking*, 28 *Law & Phil.* 367 (2009).

³⁹ Other doctrinal consequences do not perfectly coincide with the individualized/ statistical distinction. For example, insurance and worker’s compensation exclusions for “intentional” torts more often apply when the knowledge is individualized than when it is statistical, but the correspondence is not exact. See, e.g., =

⁴⁰ The point in the text also applies in tort law. Consider the examples mentioned in note 8= supra [=quote from R.3rd torts=]. If the actors in any of these cases unreasonably took inadequate precautions against harm, they would be considered negligent but would not be considered liable for the intentional tort of battery.

cases.⁴¹ Statistical knowledge cases, if culpable enough to deserve legal sanction, will ordinarily be classified as cases of recklessness or negligence, not as cases of individualized knowledge.

Second, although individualized knowledge sometimes differentiates criminal from noncriminal behavior (e.g., a crime might prohibit knowing but not reckless possession of a weapon), again statistical knowledge will not do as a substitute for such knowledge. Recall the example of a large employer charged with knowingly hiring an undocumented worker. If he utilizes the best practicable method for screening out such workers, but the scale of his operation is so large that he knows that some of his workers must be undocumented aliens, it is exceedingly unlikely that a court would permit his conviction.

And third, the demand for an especially strong justification where the defendant has individualized knowledge that he will cause harm does not ordinarily extend to defendants with statistical knowledge: the law will harshly sanction Clara, who lacks a socially compelling justification, but will not sanction Agatha at all, even though her justification is arguably not much more compelling than Clara's.

But why are statistical knowledge and individualized knowledge treated so differently? And what, exactly, is the relevant difference? The answer to these questions is vitally important, not only in differentiating more culpable actions (knowing) from less culpable ones (reckless or negligent), but also in differentiating those "knowing" actions that are culpable or deficient from those that are entirely permissible or justifiable.

A central rationale for the difference between individualized and statistical knowledge is that causing serious physical harm with individualized knowledge is almost always unjustified, and indeed almost always highly unjustified. Consider Clara again. This is a paradigm individualized knowledge case because the actor knows that there is a very high probability that she will cause serious harm to a single, identifiable person. Here, the law has good reason to distinguish knowledge from lesser forms of culpability. The decision to take an action that one realizes is very likely to seriously harm another is virtually always unjustifiable: it can only be justified by a small number of extremely narrow justifications, most prominently self-defense and necessity.⁴² All else being equal, one who knowingly creates a higher risk of harm needs a weightier justification in order to deserve exemption from legal responsibility than one who knowingly creates a much lower risk of that harm.

To be sure, the morality and social acceptability of imposing risks is a complex topic. But, whether one endorses consequentialist or deontological principles (or some stew of the two), and whatever one's views about what risks are permissible or impermissible, it is not controversial that ordinarily an actor needs a

⁴¹ Some famous tort cases seem to be striking counterexamples to this proposition, especially the Ford Pinto case; but they are not. See = *infra*.

⁴² As noted above, many states that permit a necessity defense do not permit it as a defense to murder. Note 42= *supra*. So an actor who qualifies as a murderer because he has individualized knowledge that he will cause harm cannot be justified in such a jurisdiction. There are reasons to question this limitation of the necessity defense. (If the pilot of a disabled airplane deliberately steers the plane from its current path to another direction, in order to minimize the almost certain death toll from twenty to five, should he really be guilty of murdering the five?) But putting that aside, notice that the limitation is especially objectionable if statistical knowledge counts as individualized knowledge; for then every corporate decision to market a product known to cause a small number of deaths in the population of users might qualify as unjustified murder.

much weightier reason to justify imposing a much greater risk (*ceteris paribus*). Clara's interest in arriving on time to an important business meeting is, in this context, an interest of only modest social value. The same would be true if her reason for hurrying home was to prepare dinner or give her children more help with their homework. In either the original Clara example or this variation, these interests would nevertheless justify her choosing to impose very small risks of injury on pedestrians by driving carefully rather than walking. But they do not justify her consciously imposing very substantial risks by driving incautiously or very fast.

Moreover, in a standard individualized knowledge case, the law properly treats the actor harshly—and more harshly than an actor whose conduct is similarly unjustified but who demonstrate a lesser culpability than knowledge (for example, cognitive recklessness or negligence). Ordinarily, the greater the probability of harm that the actor knowingly creates, the more the actor devalues the interests of the victim (or other socially valued interests) relative to the actor's own ends.⁴³ Knowingly taking a high risk of seriously injuring someone for the purpose of one's career is a different matter than taking a very slight risk of such an injury for the same reason.

What, then, is the appropriate culpability level for those who act with statistical knowledge? Which of the culpability categories (other than individualized knowledge) most closely corresponds to their culpability? Recklessness? Negligence? Or no culpability at all? The short answer is: any of the above could be the appropriate culpability level, depending on the context. Let us see why.

Consider two somewhat stylized statistical knowledge cases:

Larry, the frequent careful driver

Larry routinely drives carefully, very near the speed limit. He is unusually well-informed, and knows the aggregate risk of injury that this poses over his lifetime of driving. (He averages 15,000 miles per year.) Specifically, he knows that it is highly likely that he will cause: (a) damage to another's vehicle or property on at least six occasions; and (b) minor personal injury to someone else on at least one occasion.⁴⁴

⁴³ Also, we normally cannot be sure that the actor who knowingly creates a smaller risk would have been willing to knowingly pose the greater one. In the unusual situation when we can be sure of this, however, arguably we have reason to sanction the actor as if he had possessed the greater degree of awareness of risk. See note = supra.

⁴⁴ I invented these figures. But whether this is an accurate depiction of the aggregate lifetime risks of driving is not critical to the analysis.

Apparently the annual risk of death in a transportation accident is 1 in 6,000. See http://www.nsc.org/news_resources/Resources/res_stats_services/Pages/FrequentlyAskedQuestions.aspx#question10. The lifetime risk of death in an automobile accident is 1 in 84.

http://www.nsc.org/news_resources/injury_and_death_statistics/Pages/TheOddsofDyingFrom.aspx.

It is interesting to consider how an actor should be treated if the risks that he subjectively perceives turn out to be either less than or greater than the risks that a reasonable person in his shoes would perceive, or than the risks that the most informed human being could possibly perceive (either ex post or ex ante). These questions are beyond the scope of this paper. Of course, if an actor unreasonably underestimates the seriousness of the risks she is posing, then even if she poses only statistical rather than individualized risks, the law might consider her negligent.

Maxwell, the frequent speeder

Maxwell routinely drives ten miles per hour over the speed limit. (Like Larry, he averages 15,000 miles of driving per year.) He, too, is unusually well-informed, and knows that it is highly likely that he will cause: (a) damage to another's vehicle or property on at least 30 occasions; (b) minor personal injury to someone else on at least five occasions; and (c) moderate personal injury on at least one occasion.

Despite his statistical knowledge that he will, over time, cause harm to others, Larry should not be considered as culpable as an actor with individualized knowledge. Indeed, he probably should not be found legally responsible at all.⁴⁵ Maxwell, on the other hand, is a more complex case. In an important sense, he seems like Larry, insofar as he has statistical but not individualized knowledge of the harm that he will cause, and insofar as the number of times he drives does not seem to affect his level of culpability. But as we shall see, there is also some reason to conclude that his repetition of an unjustifiable action legitimately bears on his culpability.

The fundamental question concerning statistical knowledge is whether such cases are better categorized along with individualized knowledge, on the one hand, or with recklessness, negligence, or even no liability, on the other. The short answer? Statistical knowledge ordinarily belongs with recklessness (or a lesser culpability level), not with individualized knowledge. This conclusion follows when we examine the dimensions discussed above, aggravated culpability (features (1) and (2)) and clear lack of justification (feature (3)).

First, consider aggravated culpability. When a person acts despite statistical knowledge that harm will result, he need not significantly devalue the interests of the potential victims. Indeed, he might well value those interests more highly than a cognitively reckless actor who is aware only of a significant risk that the same harm will befall a particular individual. For example, compare: (1) the decision of a delivery service to use automobiles rather than bicycles in a particular municipality, despite statistical knowledge that this will over time cause more injuries to pedestrians (but will greatly speed delivery time); with (2) the decision of a motorist or bicyclist, on a single occasion, to speed very near a particular pedestrian, aware of the significant risk of injury. The second actor is clearly more culpable, even if she correctly estimates the risk of injury from her single act as 10% while the delivery company correctly estimates the risk that its (carefully operated) fleet of cars will at some point injure a pedestrian as 95%.

Second, consider the question of justification. The heavier burden of justification that we demand when an actor knowingly creates a high (rather than a low) risk to a victim on a single occasion does not imply that we should require a similarly high burden of justification when an actor's "knowledge" that others are likely to suffer harm is merely statistical. Statistical knowledge cases simply involve the predicted aggregate effect of numerous low-risk individual acts; but if a low burden of justification is appropriate in the case of each such act, then it should ordinarily be appropriate when we aggregate the acts. The decision to engage in an

⁴⁵ I qualify this assertion because imposing strict tort liability on Larry for the harm he causes, or requiring him to insure in advance against such harms, might not be unfair. I do not have space to explore the issue here.

activity that one realizes is statistically likely to cause harm (given the repetitive or far-reaching nature of the activity) carries no presumption of unjustifiability. To be sure, the activity might not be justifiable; but whether it is or not should normally turn on the usual criteria relevant to justifiable risk-creation, not on whether one can recharacterize the activity as if it were a single event stretching over time and space. Or, to put the matter differently, how one frames the scope of the risk is quite arbitrary: one might describe the risks of such activities 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 40%; but the chance that my moderate speeding will injure someone in a single trip to the beach is vanishingly small.⁴⁶ The legal permissibility of the risky conduct, and how (if at all) it should be sanctioned, should not depend on the contingency of this arbitrary risk description.

These two arguments are, I recognize, too brief to explain away the statistical knowledge conundrum. But they are a start. And we will make further progress if we proceed to a more formalized explanation of the reason we should ordinarily treat statistical knowledge differently than we treat individualized knowledge.

B. A partial solution: the principle of ICAA (Invariant culpability when acts are aggregated)

Larry, the frequent careful driver, should not be found legally at fault for property harm that he causes just because he knows, in a statistical sense, that this result is a highly likely consequence of his taking tens of thousands of drives. After all, each time he drives, the risk of property damage (and of personal injury) is extremely small. Assuming no other fault in the way he drives on any particular occasion (such as inattention or a dangerous maneuver), Larry is not negligent or otherwise at fault for choosing to drive at such a speed. This is so even though safer alternatives exist, such as walking, or riding his bicycle, or driving much more slowly. And if Larry does not act in a faulty manner on any of his individual drives, there is no apparent reason why his decision to repeat his activity thousands of times should transform the activity from faultless to faulty.

Implicit in this argument is the following principle:

⁴⁶ See Simons, *Negligence*, supra note 36= at =. The point in the text is an aspect of the notorious “reference class” problem in probability analysis. On a frequentist analysis of risk, one needs to explain what class of persons the frequencies refer to, but there is no uncontroversial way to do so. See Adler, *Against ‘Individual Risk,’* supra note =, at 1144-1147.

Or consider the infamous Ford Pinto case, where Ford was found to have negligently designed an automobile in such a way that the fuel tank was unusually likely to cause serious fires in the case of a collision. Ford conducted a cost-benefit analysis, using figures supplied from the federal government, that they claimed excused their failure to make a post-sale modification of the vehicles to prevent such fires. Reportedly, 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. That sounds like an insignificant percentage. However, we could alternatively describe that risk by the expected number of burn deaths over the entire fleet of Ford Pintos. Framed this way, the risk is an expected toll of 180 expected deaths over the fleet of Pintos. See Simons, *Rethinking Mental States*, supra note 15= at 292 n. 69; Gary Schwartz, *The Myth of the Ford Pinto*, supra note 17=, at 1059-1062.

Invariant Culpability when Acts are Aggregated (“ICAA”):

Suppose: (1) a person engages in a course of conduct containing multiple acts each of which, considered individually, is identical in all respects relevant to legal culpability, including:

the actor’s intentions, beliefs, and motives;

the socially recognized benefits and risks of harm that the act entails; and

the manner in which the benefits, and risks of harm, from the act are distributed among the actor, potential victims, and others;

Then: (2) we should attribute to that actor the same degree of culpability for the whole course of conduct as the degree of culpability that we properly attribute to him for any individual act.

Under ICAA, if performing a single act is faultless, then performing a hundred relevantly similar acts is faultless. If performing a single act is merely negligent, then so is performing a hundred similar acts; and so on for other types of culpability. Aggregation of acts does not, by itself, increase culpability. ICAA is especially crucial to understanding the proper scope of a cognitive mental state such as knowledge. For, as the scope and frequency of otherwise similar acts increases, the actual and reasonable beliefs of the actor about the probability that his aggregate acts will cause harms (or will produce benefits) will typically change: he will recognize that those harm are now much more certain. But, absent a constraint such as ICAA, the actor will implausibly be liable for a higher level of fault, or for some level of fault even though each separate act is lacking in fault, simply because he expands or repeats the activity.

ICAA presumptively should govern when similar acts are multiplied either over time or over space. If Alicia only drives once in her lifetime, while Bella drives 10,000 times, and Carmen 300,000, ICAA provides that their responsibility for resulting harm should be the same, so long as their driving behavior is the same.⁴⁷ And if a company expands its fleet of drivers ten-fold, or 100-fold, or 10,000-fold, but otherwise maintains the same safety policies, its liability or responsibility for harms caused by driving should not change, simply because the aggregate risks are much greater.⁴⁸

⁴⁷ The number of “drives,” a notion that I employ for simplicity here and in many other examples, is imprecise in several ways. For one thing, a single drive contains numerous acts, as I discuss momentarily. For another, miles driven, weather conditions, congestion, and other factors are highly relevant to the risk posed. For purposes of this essay, assume that the characteristics of a single “drive” are reasonably well-defined, so that we can meaningfully (a) compare different actors each of whom takes a single (relevantly similar) “drive,” and also (b) compare an actor engaging in one “drive” to that actor engaging in many (relevantly similar) “drives.”

⁴⁸ Of course, insofar as harms actually caused increase, compensatory liability for those harms must increase; and insofar as the number of risky acts increases, sanctions for those risky acts must proportionately increase. But ICAA provides that the level of fault attributable to an act risking or causing harm does not vary with the number of similar acts committed. Nonfaulty conduct does not become faulty. Negligent conduct does not become reckless. And reckless conduct does not become knowing (in the individualized sense).

It might seem that statistical knowledge should be treated like individualized knowledge at least in a case where it is very likely that the activity in aggregate will cause at least one instance of serious physical injury or death. But this argument is flawed, because it is vulnerable to the largely arbitrary way in which we frame our spatiotemporal description of an activity and its effects. Over a million years of repetition, innumerable human activities will cause physical harm.⁴⁹ Over a span of a second, few activities will. Yet neither description is more or less accurate than the other. Moreover, although the relative urgency of precaution depends in part on the frequency of harm over time, even a very low-level risk is worth averting if the burden of doing so is sufficiently low and if the harm (were it not averted) is sufficiently great.

The discussion above of how ICAA applies to the number of “drives” a person is inexact. To be more precise, we could try to individuate, not separate “drives,” but instead distinct but comparable segments of risky behavior, or even distinct acts of risk creation. After all, a given drive could easily contain numerous discrete episodes or acts of risk creation—for example, driving at one speed through the first intersection, at another speed through the second, steering close to the sidewalk at the third, steering farther away at the fourth.⁵⁰ Suppose Alicia drives for fifteen minutes in such a way that she poses a tiny risk to one person before she concludes the drive, while Arthur drives in such a way that he poses an equivalent risk successively to three people, endangering each of them for fifteen minutes along different parts of the road. Alicia, we might say, has engaged in one episode of risky behavior, while Arthur has engaged in three comparable episodes.⁵¹ Along each road segment, each has engaged in bodily movements, and has made a number of decisions such as how to steer his car and how fast to drive.

But perhaps we should ask the narrower question: How many distinct acts has Alicia or Arthur performed? Alas, act individuation is a notoriously difficult conceptual (and perhaps normative) problem, in both philosophy and criminal law

⁴⁹ If you were immortal, yet persisted in engaging in risky acts such as driving, or even trivially risky acts such as taking a brief stroll down a quiet street, you could be sure to cause a death eventually. But it is doubtful that you should therefore be considered a prima facie murderer. See Tyler Cowen, http://volokh.com/2004_01_11_volokh_archive.html#107409984581699780; Tyler Cowen, http://www.marginalrevolution.com/marginalrevolution/2004/01/would_potential.html, quoting a post by Lawrence Solum.

⁵⁰ A single drive might also endanger no persons, one person, or many persons. Indeed, a single act within one “drive” could differ in how many are exposed to risk, and in the type of risks of harm that it creates. This second aspect, the number of risk-exposures, needs to be analyzed separately. Obviously, a single act of risk-creation (driving near the sidewalk) could endanger numerous persons or one, or could create serious risks of death and personal injury or only minor risks of property damage. I provide such an analysis below, introducing the related but distinct principle, ICREA (Invariant culpability when risk exposures are aggregated). *Infra* notes =.

⁵¹ Further complexity is added if we apply ICAA to acts of risk-creation in which no one is “objectively” at risk (say, from the perspective of a witness at the scene), e.g. if Arthur drives at night in a college town and no pedestrians are nearby. Should we examine the risk according to Arthur’s subjective beliefs? According to the risk that a reasonable person in Arthur’s shoes would perceive? According to the risk that an omniscient observer would perceive? How we should characterize risk is a notoriously difficult problem in criminal law and moral theory. I believe, however, that the proper resolution of these problems does not affect the arguments I make in this paper. A range of plausible criteria of “risk” are consistent with the basic argument that ICAA should presumptively apply if an actor engages, not just in one act of “risk”-creation, but in multiple similar acts.

theory.⁵² Fortunately, for purposes of our inquiry, we can safely skirt the problem. Whatever conception of “act,” or of risky episode, we apply to Alicia, we can readily construct scenarios in which another actor (Arthur) engages in several acts or episodes of that nature. The question then remains: is Arthur, by reason of repeating such acts, more culpable than Alicia? And ICAA gives the presumptive answer: no. I will also employ the idea of a “single drive” as a crude proxy for a single act while driving that creates a risk of harm. This idea permits a more straightforward comparison of different scenarios and is thus more useful in identifying the strengths and limitations of ICAA.⁵³

What deeper theory of permissible risk-creation and of punishment underlies the ICAA? Perhaps the most straightforward defense of ICAA is provided by a consequentialist⁵⁴ analysis under which the justification for engaging in an activity with a known set of risks depends on the ex ante expected costs and benefits of the activity. Thus, under the Learned Hand test,⁵⁵ if it is not negligent for Larry to drive at a moderate speed rather than walk, this means that the burden of not driving on a given occasion is greater than the marginal costs of driving on that occasion (where risks of injuring others are prominent among these costs). But multiplying occasions simply amounts to multiplying each side of this equation by the same number. The resulting equation will obviously preserve the original inequality. If, for one drive, $B > P \times L$, so that the actor is not negligent, then obviously, for 10,000 drives, $[10,000 \times B] > [10,000 \times (P \times L)]$, and the actor is still not negligent.⁵⁶

⁵² See, e.g., Michael Moore, *Act and Crime* (1993); Alexander & Ferzan, *supra* note =, at =.

⁵³ One further reason justifies agnosticism about how to individuate “acts.” The only concept that we really need in order to defend ICAA and to explain why statistical knowledge is distinguishable from individualized knowledge is the concept of a shorter rather than longer episode of conduct in which the actor (1) creates risks of harm, (2) is potentially culpable, and in which (3) the episode is repeatable. The episode could, but need not, be an “act.” Let me explain.

Individualized knowledge, as I define it, means that an actor believes he is imposing a risk of harm on the particular victim greater than some threshold probability T (e.g., 75%). (See Appendix.) So if an actor believes that the risk of harm to the victim due to his acts or conduct during the relevant episode is as large as T , then he causes that harm with individualized knowledge. If, however, he believes that the risk is less than T , then he lacks individualized knowledge. The point of ICAA is to emphasize the following: if $A1$ engages in a short episode of conduct that does not create a risk to an individual victim as high as T , while $A2$ engages in a longer episode that also falls short of creating a risk to an individual victim as high as T , then: (a) neither has acted with individualized knowledge, and (b) $A2$'s culpability is presumptively no greater than $A1$'s.

I thank Kim Ferzan for pressing me on the “act individuation” problem, though I am confident that she will find this resolution unsatisfactory.

⁵⁴ This form of consequentialism, however, is truncated: it examines only whether a particular type of act or behavior is optimal or instead deficient or unjustified. A broader consequentialist account will also consider whether it is optimal to impose legal sanctions on behavior that is suboptimal in the narrower, truncated sense. See Simons, *Tort Negligence*, *supra* note 13=, at =.

⁵⁵ The Learned Hand test need not be understood in consequentialist terms, however. See Simons, *The Hand Formula*, *supra* note =, at =; Simons, *Tort Negligence*, *supra* note 13=, at =.

⁵⁶ See Richard A. Posner, *Economic Analysis of Law* 261 (8th ed. 2011):

Consider a railroad that because it runs many trains every year knows with a confidence approaching certainty that it will kill 20 people a year at railroad crossings. Is it therefore an intentional tortfeasor? It is not, either in law or in economics. The same thing that makes *PL* high—the scale of the railroad's

But a nonconsequentialist also can, and indeed should, endorse ICAA. Whatever criterion of unjustifiable risk-creation a nonconsequentialist endorses for a single act, the criterion should apply *mutatis mutandis* to multiple acts of the same type. For example, suppose the criterion is this: it is only justifiable to impose a risk of serious injury on others so long as those others either (a) consent to or benefit from the risk or (b) impose a risk of similar magnitude on the actor. Then if a single risky act by Larry is sufficiently reciprocal in the risks it creates to satisfy this criterion of permissible risk-imposition, multiple acts of that type are also sufficiently reciprocal to satisfy the criterion. And the analysis should be the same for most other plausible nonconsequentialist criteria.⁵⁷

Consider a highly abstract depiction of the argument:

Chart 1: Larry

	Disvalue of risk (of causing personal injury)	Value of benefit
Larry, one drive [Not at fault]	1	2
Larry, 20,000 drives [Not at fault]	20,000 [aggregate 70% risk of causing injury]	40,000

The example is stylized insofar as the numbers represent magnitudes of values that cannot really be quantified with any precision.⁵⁸ Thus, I am simply positing that

operations—makes *B* high. The ratio of *B* to *PL* is unaffected by the scale of the potential injurer’s operation, and it is the ratio that enables us to differentiate between intentional and unintentional torts in an economically relevant sense.

⁵⁷ For example, drivers on a highway or bicyclists on a bicycle path pose significant risks of harm to each other. If those risks are sufficiently reciprocal to be justified on a nonconsequentialist account, then it is justifiable for a driver or bicyclist to impose them once, twice, or a million times. (For some nonconsequentialist accounts of unjustifiable risk-creation, see note 36= supra.)

At the same time, I concede that ICAA will not hold for nonconsequentialist criteria that include thresholds. And sometimes, as I suggest below, thresholds are plausible to include (as when environmental risks are harmless below an aggregate threshold and harmful above). Often, however, thresholds are very problematic. Nonconsequentialist views that very broadly prohibit the imposition of “substantial” or “significant” risks have to be formulated very carefully in order to avoid the arbitrary framing objection I have noted. Cf. Keating =; Wright =.

⁵⁸ Moreover, for purposes of both economic analysis and other approaches (including nonconsequentialist principles) that analyze the justifiability of risk-creation, one should focus on the marginal risks and benefits of a particular level of care. See Simons, Tort Negligence, supra note 13 =. For simplicity, I do not emphasize this marginal feature in the charts, but it would be possible to modify the charts to reflect it.

“20,000” is an abstract measure of the disvalue of creating a 70% risk of causing personal injury. And I am not assuming a consequentialist justification for the legal sanction or regulation of the risky conduct, though I am indeed assuming that the justification (whether consequentialist or nonconsequentialist) requires some type of balancing of socially recognized risk and socially recognized benefit.⁵⁹ But the chart is still useful in revealing significant relationships (as will become more apparent shortly when additional examples are added).

Chart 1 illustrates one important lesson of ICAA: if we initially and properly characterize behavior as not faulty, we should not deem the behavior faulty simply because it occurs more often. Suppose a known 70% probability of causing personal injury (in a case such as Clara or Eleanor) suffices as individualized knowledge. ICAA provides that if Larry drives carefully, then even if over 20,000 drives he knows there is a 70% probability that his driving will cause personal injury to someone, he should not be treated the same as those actors possessing individualized knowledge. We should not elevate his culpability simply because he drives more frequently.

To be sure, ICAA sometimes needs to be qualified in order to address exceptional situations. One exception involves thresholds of harm. If an actor contributes to a social harm, and that harm is significant only if it exceeds some threshold, then the actor might owe a special duty not to contribute to crossing that threshold. Some environmental harms fit this description: certain contaminants are hazardous to health only if the aggregate air or water pollution that they create exceeds a particular level. In such circumstances, ICAA should not apply strictly. The fact that a farmer or manufacturer could justifiably produce a small amount of pollution in light of the social benefits of his activity does not mean that it can justifiably produce a much larger amount of pollution, even if the social benefits are also proportionately greater.⁶⁰ Similarly, a driver might not be justified in taking as many trips as he likes in a city where aggregate air pollution from automobiles is a health hazard, even if he would be justified in taking that many trips in an unpolluted city. However, these types of cases are exceptional, and ICAA can be qualified to accommodate them. By contrast, for the kind of risky conduct and activities most commonly at issue in tort and criminal law, there is no threshold which the actor has a special duty not to exceed, nor is there an invariant threshold (of probability or of expected harm) below which the risk of harm is inconsequential. Aggregate property damage and physical injury are social harms that normally increase in rough proportion to increases in such risky activities as

Furthermore, the permissibility of an act and the degree of its reprehensibility cannot be adequately expressed as simple functions of two continuously varying values, at least on any moral or legal theory more subtle than economic cost-benefit analysis. The number of relevant values could be much more numerous, and the way in which they matter—the structure of wrongdoing—could be much more complex. For one recent, promising effort to formalize nonconsequentialist constraints on consequentialist balancing, see Eyal Zamir & Barak Medina, *Law, Economics, and Morality* (2010).

⁵⁹ See Simons, *Tort Negligence*, supra note 13=.

⁶⁰ For discussion, see Joel Feinberg, *The Moral Limits of the Criminal Law, Vol. I: Harm to Others* (Oxford 1984), pp. 225-232.

A range of scenarios needs to be addressed in order to fully analyze the threshold of harm issue. For example, even a tiny contribution from one manufacturer might, when combined with tiny contributions from 200 other manufacturers, exceed the threshold, in which case all of the manufacturers arguably have a duty not to contribute to the harm.

driving, distributing a particular type of product, constructing a particular type of building, or performing a particular type of medical procedure.

C. Does ICAA apply to faulty actors?

Now consider whether ICAA applies to faulty actors. Does it preserve whatever level of fault the actor displays in a single action, even if the actor repeats the action? Or can repetition increase the level of fault?⁶¹ Recall the example of Maxwell, who routinely chooses to drive ten miles per hour over the speed limit. In any individual drive, he is negligent and at worst reckless.⁶² If he drives a thousand times, does he become a knowing wrongdoer? If his speeding causes a death on his nine hundredth drive, is he guilty of murder (which typically is defined to include knowing causation of death) or instead of manslaughter (often defined as reckless homicide)? On first glance, the number of times that he engages in a particular type of act does not seem to change his culpability; the ICAA principle seems to apply. Thus, manslaughter seems to be the more appropriate level of punishment. Of course, the fact that he repeatedly acts unjustifiably certainly warrants imposing aggregate legal sanctions (whether in the form of civil fines, civil liability, or criminal punishment) that reflect the aggregate risks and harms he has created. But ICAA instructs that the warranted total sanction is merely additive, equaling only (but no more than) the sum of the warranted sanctions for each act.⁶³

Yet this initial conclusion, that ICAA should extend to the benefit of faulty actors, may be too hasty. In the earlier example, Larry took a justifiable risk, and ICAA properly applied. But perhaps ICAA does not apply when an actor such as Maxwell takes unjustifiable risks (or otherwise demonstrates culpability) in a series of acts. Here, it seems, the whole can be more than the sum of its parts. On this view, Maxwell's speeding twenty times is a worse state of affairs or reflects greater aggregate fault than an alternative state of affairs in which twenty different, unrelated individuals each speeds once.

⁶¹ This issue is particularly acute when the "aggregation" or multiplication occurs over time, rather than over space. The examples discussed in this section involve the former.

Most of the analysis in this section, including the argument that the culpability for the "whole" course of conduct sometimes can be greater than the culpability for a "part," does not apply when the multiplication is merely "spatial." If Domino's Pizza adopts a policy today requiring its delivery drivers to speed in order to assure prompt delivery, that negligent or reckless policy remain negligent or reckless, whether it applies the policy to ten, a thousand, or 100,000 drivers. But if it continues the policy over time, new arguments emerge for attributing to Domino's a higher level of culpability. See TAN = infra.

⁶² Is Maxwell reckless? Is he sufficiently "aware" of the risk of harm he is posing simply because he has a general, background appreciation of the risks of speeding? On some modern formulations, such as the Model Penal Code's, the state probably must show more: that he possess a more vivid or specific awareness of the particular risk that his speeding poses on the given occasion. See Simons, *Should the Model Penal Code's Mens Rea Provisions Be Amended?*, supra note 15=, at =.

⁶³ I do not mean to suggest that a retributivist must endorse a simple additive formula for converting desert into punishment. Even if D1's criminal desert is twice that of D2, it hardly follows that D1 ought to be incarcerated for precisely twice as long as D2. How one converts retributive desert into actual institutional punishments raises difficult questions of quality and quantity of punishment that I do not explore here. Similar difficulties attend the task of converting an actor's culpability into tort punitive damages on a corrective justice or other nonconsequentialist tort principle. If D3 is twice as culpable as D4, it does not follow that the punitive award that D3 should pay ought to be twice as large as the award that D4 should pay.

Consider the fact that the criminal laws of most states punish recidivists especially harshly.⁶⁴ In this respect, they depart from ICAA, for they impose on the recidivist a punishment significantly greater than would be warranted if he had committed only his most recent offense and no prior offenses. (Put differently, his total punishment for all offenses will exceed—often by an enormous amount—the sum of the distinct punishments that would otherwise be warranted if each of the recidivist’s offenses had been committed by different individuals.⁶⁵) Similarly, a tort defendant who repeatedly endangered others (e.g. someone with a long history of drunk driving) would be a prime candidate for punitive as well as compensatory damages, even if the dangerous acts, if committed by distinct individuals, would only warrant compensatory damages.⁶⁶

A consequentialist has good reason not to apply ICAA to (individually) faulty actions. If Maxwell repeatedly acts unjustifiably, a greater than usual sanction will often be needed in order to specifically deter him, since he (by hypothesis) has not been deterred by the normal threatened sanction for the prior acts; and similarly, general deterrence might be undermined unless we imposed a greater sanction for actors who repeatedly commit wrongs. Thus, consequentialist analysis does justify recidivist sanctions (and comparable sanctions) that depart from ICAA. However, for Larry, who repeatedly acts justifiably, consequentialism requires adherence to ICAA. A cost-benefit analysis suggests no reason for recharacterizing his culpability as faulty rather than faultless, simply because he repeats his acts. After all, by hypothesis, the net social value of a single act (judged

⁶⁴ To be sure, it is controversial whether and why we are justified in imposing greater civil or criminal sanctions on recidivists and those with a record of drunk driving than ICAA would permit. But the justification, if any, cannot be simply that recidivists realize that harm is proportionately more likely to occur if they commit risky activity more frequently. If that were the only reason, there would be no justification for enhanced punishment of recidivists. Other possible reasons are examined in the text below.

⁶⁵ But note that often the law takes the opposite tack and grants an effective “bulk discount” to a different type of “recidivist,” the offender who commits numerous crimes prior to being caught. The state will typically punish Maxwell much less for speeding twenty times before being apprehended than the aggregate punishment that it will impose on twenty unrelated individuals who each independently speeds once. How (or whether) retributivists can explain such a discount is a difficult question. See Jesper Ryberg, *Retributivism and Multiple Offending*, 11 *Res Publica* 213 (2005); Alexander & Ferzan on “volume discounts” in punishment. =. Perhaps, when one has committed not two or three but fifty burglaries, a discount is appropriate because the behavior expresses a compulsiveness that is somewhat mitigating.

However, there are many complicating factors at play here, beyond principles of just deserts. For example, a judge might give the one-time speeder a higher punishment than his conduct by itself deserves in order to teach him a lesson, or because she believes he has likely engaged in numerous similar risky acts for which he escaped apprehension or sanction.

⁶⁶ One might further distinguish between “de jure” recidivists, who have been convicted of a prior crime or found liable for a prior tort, and “de facto” recidivists, who have engaged in similar criminal or tortious acts in the past but have not been legally sanctioned for them. (See prior footnote.) The extra culpability of a de facto recidivist arises from his persistence in endangering others. His defiance of his legal obligations is at most implicit. The extra culpability of a de jure recidivist is arguably greater, arising both from such persistence and from his explicit defiance of his legal obligations. = See TAN = next page.=

ex ante) is positive, so the net social value of a series of identical acts will also be positive and proportionately greater.⁶⁷

But on a nonconsequentialist account, does ICAA continue to apply even to wrongful and unjustifiable conduct? The answer is less straightforward. One important complication here flows from an empirical observation. When someone repeats such conduct, he often does so with a higher level of culpability than before—for he might, in the process of repeating the act, either develop a new level of awareness of the risks, or act with an especially egregious motive or attitude not present in a single culpable act. However, on closer inspection, such cases are not really a counterexample to ICAA; rather, they are simply instances in which, as an empirical matter, ICAA's ceteris paribus condition does not apply. Let me explain.

As noted above, ICAA holds true if we keep constant all other factors relevant to culpability, including beliefs, intentions, motives, benefits, and harms. But as an empirical matter, the repetition of wrongful acts often will be accompanied not only by mere repetition of these culpability factors, but also by new elements or degrees of culpability. Thus, in practice, an actor (such as a speeding driver) will often have a more confident and certain belief about the harmful results of his conduct as he repeats the conduct, especially if the conduct occasionally results in harm.⁶⁸ For example, even if a single incident would only be judged negligent, repetition of the very same conduct will often be judged reckless, simply because the evidence is now clear that the repeating actor has become cognizant of the risks.⁶⁹ And similarly, when the conduct repeats, a legal decision-maker will also have more confidence in judging that the actor was aware of the risk all along. Moreover, as the number of acts multiplies and the time frame expands, the actor can no longer rely on the claim that the act was impulsive and thus less culpable.⁷⁰ And since he

⁶⁷ See Posner, *supra* note =. This analysis requires qualification. It is possible that the single act, considered alone, does not have positive net social value, but is not worth regulating by legal sanctions because the costs of such regulation are excessive. Yet it is quite conceivable that a series of similar acts are worth regulating, if the marginal costs of regulation decline sufficiently with the increased number of acts. (The state might not find it worthwhile to prosecute a petty theft unless the offender is known to have committed numerous small thefts.)

⁶⁸ On the other hand, if an actor repeatedly engages in risky conduct yet fortuitously does not cause harm, he might lower his subjective probability estimate of the risk. On a purely subjective understanding of culpability, this will lower his culpability. But it is often unreasonable to respond to nonoccurrence of harm by lowering one's subjective estimate of the ex ante risk. Speeding drivers only rarely cause harm, but they can become irrationally overconfident that their speeding is not risky until they actually cause an accident. One solution to this problem is to permit legal liability for running risks that are wrongful in some "objective" sense (e.g. from the perspective of a reasonable person).

⁶⁹ Indeed, faultless conduct can become faulty due to these effects of repetition. An actor might not realize, and might be blameless for not realizing, that a particular risk exists in an unfamiliar and rare context—for example, that it is worthwhile to check the brakes of a rental car before leaving the rental lot. But if he rents cars every month and on one occasion discovers this problem, perhaps he has a duty to check thereafter.

⁷⁰ With impulsive, sudden reactions, even if the actor is aware of the harm his action might cause, he is sometimes less blameworthy because he merely has an intellectual, not affective, awareness of the significance of that result, or because he has not really considered the result, much less considered it carefully. See also Robert Nozick, *Philosophical Explanations* 392 ("[R]epetition [of a criminal act] shows that more of the person, a deeper aspect, is behind it—it was not just one isolated passing whim.").

Similarly, if the actor's knowledge is latent, rather than uppermost in his consciousness, that might be a reason to mitigate his legal liability. See, e.g., Moore & Hurd, draft on

has more opportunities for reconsideration of his acts and more opportunities to take mitigating precautions, his failure to take such steps suggests a more considered elevation of his own interests over the welfare of others, and thus adds to his culpability. Indeed, certain actors (such as corporations) sometimes have an affirmative duty to disclose past wrongs; continued nondisclosure might then elevate the wrongdoing to the level of fraudulent concealment.

Nevertheless, I do not believe that this empirical gloss is the entire (nonconsequentialist) story. Apart from these empirical contingencies, both the recidivist premium and similar policies of aggravation seem defensible in principle. Specifically, it seems justifiable:

- (1) to punish recidivists more harshly than they would deserve if they had only committed their most recent crime;
- (2) to impose especially severe punishment on actors who successively create multiple risks of harm in a single extended criminal episode;⁷¹ and
- (3) to impose punitive damages on corporations that continue to market a dangerous product even after they realize the danger.

These positions rest on a rejection of ICAA as normatively unattractive in this context, even if it might sometimes be empirically applicable. An actor who persists in wrongdoing repeatedly is legitimately viewed as committing a wrong more serious than the composite wrong that is constituted by the sum of the individual unjustifiable acts. And this is not merely because, empirically, the course of conduct is likely either to produce higher levels of awareness or to provide more opportunities for reconsideration that we can fault the actor for not exploiting.

Rather, I believe the explanation is this: such persistence often reveals the culpability of defiance. This defiance of legal and social norms is a distinctive type of wrong, a wrong that transcends the culpability of the individual acts, and might properly warrant enhanced legal sanction.⁷² The whole can indeed be greater than the sum of the parts.

The issue remains, however: just how culpable are such “repeat offenders”? Specifically, we should ask two critical questions about actors who repeatedly engage in culpable acts:

- (1) Does the statistical knowledge that such actors possess genuinely amount to enhanced culpability?
- (2) If so, does it boost culpability all the way up to the level of individualized knowledge?

negligence=. But as the actor continues to engage in numerous acts of a similar nature, his knowledge of the risks that they pose will often become more “occurrent” and less latent.

⁷¹ E.g. =, where D was found guilty of depraved heart murder because of his willingness to expose numerous victims to lethal harm over an extended period of highly dangerous driving.

⁷² I do not have space here for a full discussion of the nature of defiance as a distinct form of culpability. Some valuable discussions include: Jean Hampton, *Mens Rea = ; =*. Hampton characterizes the most serious types of criminal culpability as exhibiting defiance and therefore justifiably incurring serious punishment. Her argument, however, applies to single as well as multiple acts, and thus is considerably broader than the argument I make here. See also Stuart Green, *Lying, Cheating, and Stealing: A Moral Theory of White Collar Crime* (Oxford 2006)(=); Youngjae Lee, *Recidivism as Omission: A Relational Account*, 87 *Texas Law Review* 571 (2009) (=Give description).

Those who entirely reject the distinction between statistical knowledge and individualized knowledge (at least for culpable actors) answer yes to both questions. But careful reflection supports much more qualified answers—namely, “sometimes yes,” to question (1), and “only very rarely,” to question (2). In answering these questions, we need to consider them from both the empirical and the normative dimensions just noted.

On the empirical dimension, the answer to these questions is highly contextual. Someone who speeds on the highway repeatedly might not have any greater awareness of the risks he is posing the fiftieth time he speeds than the first time, and the repetition need not reveal an especially blameworthy reason for his conduct. On the other hand, a corporate defendant that distributes a defective, harmful product might legitimately doubt the significance or even the reality of the problem if only a few incidents have been brought to its attention, but as the incidents greatly multiply, this attitude is no longer credible or defensible. And even if the multiplication of incidents merely reflects a proportionate increase in units of the product distributed over time, the fact that the corporation has a more extended time period for identifying and correcting the problem justifies treating the “repeat offender” corporation as more responsible if it fails to act.

With respect to question (2), again the answer depends on the context. Someone who seriously endangers others repeatedly during a single episode of drunk driving and then finally causes a pedestrian’s death might properly be treated as harshly as a driver who knowingly causes another’s death; indeed, the doctrine of extreme indifference murder would, in many states, treat them the same.⁷³ One rationale for this result is that such an actor ordinarily becomes increasingly aware of the risks he is imposing and of the need to take compensating precautions. Another rationale is that such an actor’s conduct is an extraordinary departure from reasonable care, so extraordinary that it is comparable to the decision by an actor (such as Clara, above) to knowingly endanger a victim on a single occasion. On the other hand, if the culpability of each of the multiple acts in question just barely satisfies the legal criteria of recklessness, then it would usually not be a proportionate sanction to treat the actor as severely as he would deserve if he had possessed individualized knowledge. One who drives just above the speed limit every day of his adult life, and thus can be fairly certain that he will cause property damage to another’s automobile on some indeterminate occasion, should not, on that occasion, be treated as harshly as one who has, on one particular occasion, driven in a highly dangerous manner and thus knowingly damaged another’s automobile.

On the normative dimension, we must consider how much normative weight should attach to the actor’s defiance (apart from the culpability expressed in the individual acts taken separately). Sometimes it is indeed appropriate to “bump up” the actor’s culpability level for legal purposes when he repeatedly engages in what would otherwise be merely reckless behavior.⁷⁴ But I believe that this aggravation

⁷³ =cites.

⁷⁴ This analysis might also apply to “negligent” behavior insofar as negligence includes not only inadvertence but also instances where the actor has a degree of awareness that falls short of what legal recklessness would require, e.g. awareness of what the actor believes to be an insignificant risk.

How does the individualized/ statistical knowledge debate apply to those negligent actors who are completely unaware of a relevant risk, but who should have been aware? Does ICAA apply to negligently inadvertent actors? Perhaps not. Perhaps as the number of risk impositions increases, it becomes more fair to blame an actor for not recognizing the significance of those risks.

in culpability is rarely enough to justify treating the actor as severely as a person with individualized knowledge. To be sure, it is usually justifiable to punish wrongdoers much more harshly when they intentionally injure or endanger others repeatedly (for example, when they repeatedly rape, murder, or rob). Yet it is quite another matter to punish a wrongdoer much more severely simply because of their more extensive knowledge of the harms they will cause or have caused—i.e., simply because, in performing a single act, the actor is aware of a risk that the law deems unjustifiable, while in repeatedly performing acts of that type, he is aware of a proportionately greater cumulative risk. The “defiance supplement,” if you will, is justifiably much greater in the first type of case (where the actor intends to harm or endanger) than in the second (where the actor displays only cognitive fault, i.e. an awareness of an unjustifiable risk).

To make this more concrete, reconsider Maxwell, who regularly speeds a little above the speed limit. Suppose his awareness of the risk on any given drive is just barely enough (in terms of the specificity of his consciousness, and the magnitude of the risk he believes he is running) to make him legally reckless. Now suppose that he takes 4,000 drives, rather than one, and suppose he is aware of the statistical high probability (say, 70%) that in the course of all of these drives, he will cause some personal injury to someone at some point. Does his statistical knowledge warrant treating him the same as Ned, who speeds through a very busy intersection on a single occasion, aware that he is 70% certain to cause personal injury to a pedestrian? Is Maxwell’s acting with statistical knowledge really as culpable as Ned’s acting with individual knowledge? Or suppose, in a variation, that Maxwell realizes that his 4,000 drives have a 70% chance of causing another’s death. After numerous drives, he does kill a pedestrian. Compare Ned taking a single highly dangerous drive, e.g. a drive in which he knowingly runs over a pedestrian while recognizing a 70% chance that this will kill her. In most jurisdictions, it is clear that Ned will be guilty of murder, while Maxwell is more likely to be guilty, not of murder, but of numerous acts of speeding plus one act of manslaughter.⁷⁵

A primary explanation for treating statistical knowledge less harshly than individualized knowledge here is that acting with statistical knowledge, while unjustifiable in this scenario, is much less unjustifiable than acting with individualized knowledge. For the individualized knowledge case is characterized by a much greater disproportion between risk and (conceivable justifying) benefit. Assume that the actor’s reason for taking the risk is the same in both cases—to

⁷⁵ One significant complication is how this analysis is affected by our rejection or acceptance of moral luck, i.e., of the idea that the occurrence or nonoccurrence of a harmful result should matter to culpability. If we reject moral luck, then Maxwell should be punished equally for all his acts of endangerment, whether or not they cause injury (so long as they are sufficiently close to causing harm that they can be considered, in effect, “last act” attempts), while Ned should be punished for his single act of attempted murder. Given the multitude of Maxwell’s acts, his aggregate punishment might exceed Ned’s.

However, if we accept moral luck, then both Maxwell and Ned receives a more serious punishment if their act results in death, and each act has about the same chance of doing so. Thus, it seems likely that accepting moral luck will tend to result in a lesser disparity of treatment between Maxwell and Ned than will rejecting moral luck.

Of course, moral luck does count in tort law, where liability for compensatory damages depends on actually causing harm.

arrive home earlier in order to relieve a babysitter and care for a sick child.⁷⁶ Again in highly stylized form, the two cases (and the case of Larry, the faultless driver) can be analyzed as follows.

Chart 2: Larry, Maxwell, Ned

	Disvalue of risk (of causing personal injury)	Value of benefit
Larry, one drive [Not at fault]	1	2
Larry, 20,000 drives [Not at fault]	20,000 [aggregate 70% risk of causing injury]	40,000
Maxwell, one drive [Slight fault]	5	2
Maxwell, 4,000 drives [Slight fault]	20,000 [aggregate 70% risk of causing injury]	8,000
Ned, one drive [Serious fault]	20,000 [single act, 70% risk of causing injury]	2

This chart reveals a dramatic difference between Maxwell and Ned: Maxwell's conduct produces a much larger social benefit. This difference legitimately bears on the appropriate legal sanction for that conduct.

When one looks at these charts, it is easy to see that Larry should not be considered legally at fault in either scenario—or more precisely, he should not be considered culpable in the second if he is not treated as culpable in the first. It is also evident that Maxwell's 4,000 drives are, in one important way, much less culpable than Ned's drive. Of course, in another way—the defiance exhibited by the repeated faulty behavior—his 4,000 drives are more culpable.

D. The special stringency principle (SSP) for highly concentrated risks

What deeper principle explains the presumptive unjustifiability of acting with individualized knowledge? A partial answer is the following deontological principle, which I will call the “special stringency principle” (SSP): one owes an especially stringent duty not to impose highly concentrated risks on another. For

⁷⁶ The structure of the analysis does not significantly change if the “benefit” is of little or no social value, such as catching the beginning of a televised basketball game, or racing a friend to one's house, though the “benefit” numbers would be much lower.

this reason, the moral and legal constraint against knowingly causing a serious harm is especially stringent, indeed disproportionately more stringent than the constraint against recklessly causing a serious harm.⁷⁷

Two more characters, Alfa and Benna, illustrate the point.

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% 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.)⁷⁸

Many consequentialists would treat Alfa-situations and Benna-situations the same, because each scenario produces the same expected net saving of lives (the "same" in the sense that each produces an identical 5:1 ratio of [expected lives saved: expected lives lost]). But many nonconsequentialist will treat the two scenarios differently in the way suggested in the quoted passage.⁷⁹

We can represent this "special stringency principle" (SSP) in graphic form:

⁷⁷ See Kenneth W. Simons, *Negligence*, supra note 36=, at 65, from which the following discussion is taken. I owe the Alfa/Benna example that follows to Leo Katz.

⁷⁸ Id.

⁷⁹ See, e.g., Kamm, *Intricate Ethics*, supra note 19=, at =; T.M. Scanlon, *What We Owe to Each Other* = (Harvard 1998). Nonconsequentialists might also consider the risk imposition more justifiable as the absolute level of the benefits and the risks decreases further while the ratio of benefit to risk remains 5:1. Suppose speeding driver Benna^ imposes a 2% risk of injury on a pedestrian in order to secure a 10% chance of saving the life of a passenger, and Benna^^ imposes a 0.2% risk of injury in order to secure a 1.0% chance of saving a passenger's life. Intuitively, the action of Benna^ is easier to justify than the action of Benna, and the action of Benna^^ is easiest to justify.

Chart 3: Alfa, Benna

	Disvalue of risk (of causing death)	Value of benefit (saving life)
Alfa, one drive [probably at fault]	1,000,000 [100% risk of causing one death]	5,000,000 [100% chance of saving <u>five</u> lives]
Benna, one drive [not at fault?]	200,000 [20% risk of causing one death]	1,000,000 [100% chance of saving <u>one</u> life]

Again, the example is stylized,⁸⁰ but it illustrates the basic point: on a nonconsequentialist approach, even if Benna's conduct is justified, more than a proportionate increase in the expected benefit is required in order to justify Alfa's conduct. Indeed, the example also illustrates a related point: knowingly causing harm, in the individualized sense of knowledge, might be impermissible even if the expected benefits outweigh the expected risks as measured by a simple risk-benefit test. This related point will be important in our discussion of corporate risk-creation, below.

But now consider this variation: suppose Benna engages in repeated acts with a cumulative risk as great as Alfa's single act? Examine the last row below, a scenario in which Benna takes five drives:

⁸⁰ The chart assumes that the benefit of saving a life is precisely equal in value to the harm of causing one death, but nothing turns on this assumption. For example, one might believe that the claim of a bystander not to be killed is significantly stronger than the claim of a dying person to be saved; and of course many other factors might be relevant to the strength of these claims. But even on these different assumptions, Alfa's conduct remains more difficult to justify than Benna's.

Also note that the benefit secured by Benna in one drive is described as "one life." This could literally mean one life, i.e., one passenger will be saved. Or it could refer to expected value, e.g., five sick passengers are in the car, and there is a 20% chance that all will certainly be saved (or it could refer to some other permutation with an expected benefit of one life). For purposes of the current discussion, these differences do not matter, though for other purposes, they might.=

Finally, I have translated "almost certain" risk of death as 100%, to simplify the exposition; but in the real world, the percentage would undoubtedly be lower—say, 90% or 95%. The same is true of the probability that the action will save lives, which undoubtedly is less than 100%. These variations do not affect the basis analysis.

Chart 4: Alfa, Benna, Repeat Benna

	Disvalue of risk (of causing death)	Value of benefit (saving life)
Alfa, one drive [probably at fault]	1,000,000 [100% risk of causing one death]	5,000,000 [100% chance of saving <u>five</u> lives]
Benna, one drive [not at fault?]	200,000 [20% risk of causing one death]	1,000,000 [100% chance of saving one life]
“Repeat Benna,” or R-Benna, five drives [not at fault?]	1,000,000 [Almost 100% aggregate risk of causing one death]	5,000,000 [100% chance of saving <u>five</u> lives]

Is the culpability of “repeat Benna,” who takes five equally risky drives, closer to that of Alfa or closer to that of “one-drive Benna”?⁸¹

The discussion earlier provides an answer. If “one drive” Benna acts permissibly, then ICAA provides that “repeat” Benna (or R-Benna) acts permissibly as well. If “one drive” Benna acts impermissibly, then R-Benna’s culpability is no greater, except to the extent that greater fault attaches to her repeated acts because they reflect her defiance of moral and legal duties. But we should not treat R-Benna as harshly as Alfa—with one important exception that “proves (the rationale for) the rule.”⁸²

Here is the exception. Suppose (and I realize that this is fanciful) that the multiple drives of R-Benna all endanger one and only one victim. For example, imagine that the route to the hospital bypasses a private home where a young child is always lingering right by the road. The 20% risk of death that this R-Benna (call her R-Benna*) poses on each drive is a risk only to this child. When R-Benna*

⁸¹ A further complication is the computation of the aggregate risks and benefits Benna creates over five drives. If the probability of harm and of benefit from one drive is independent of the probability from each other drive, which I have implicitly assumed, then we cannot simply add the risks or probabilities. Instead, the risk that Benna will cause death after five separate drives is actually $1-(1-.2)^5 = 1-(.8)^5 = 0.67$ or 67%.

This complication is ultimately not that important: we can make it irrelevant if we change the example so that the risk of death from Alfa’s one drive is 67%; then five drives by Benna creates the same aggregate risk. (Or if Alfa is characterized as causing a higher risk of death, such as 90%, then it turns out that Benna would have to take a very large number of drives, each creating a risk of 20%, in order to reach that cumulative risk—more than ten drives, in fact!)

⁸² As a contingent matter, it is also possible that as Benna increases the number of her drives, (a) she develops a clearer awareness of the risks that she is running on any given drive, or (b) in some other way her fifth drive is no longer relevantly identical to her first. See TAN = supra.

continues to drive past the child a sufficient number of times that the aggregate risk to that child is the same as the risk posed by Alfa to the victim on Alfa's single drive, there is no longer much distinction between R-Benna* and Alfa. The reason is clear: R-Benna* is now imposing all of the risk of her activity on a single person. But that is precisely why Alfa's conduct is especially difficult to justify. That is the force of the special stringency principle.⁸³

The R-Benna* example vividly highlights a crucial but underappreciated reason why an actor such as Alfa is especially culpable and is ordinarily much more culpable than either single-trip or multiple-trip Benna: an actor commits a distinctive and especially serious moral and legal wrong when the risks from her conduct are highly concentrated and distributed upon a single victim.⁸⁴ Note that this distribution-of-risk explanation is not a counterexample to ICAA; rather, ICAA simply does not apply when acts significantly differ in how they distribute the risk of harm (as I explained earlier in defining ICAA). If R-Benna's five drives endanger different classes of victims, ICAA applies: her multiple drives are no more culpable than a single drive. But if her five drives all endanger the same victim, ICAA no longer applies: that victim is now subjected to a much more concentrated risk.⁸⁵

⁸³ For a similar example, see Alexander and Ferzan, *supra* note =, at 65:

Suppose ... that the risks of tunnel building were concentrated on one known individual—Sam. Sam lives near the construction site, has a rare medical condition such that repetitive jack-hammering will eventually cause him to die, and cannot be moved. If tunnel building's benefits justify the loss of several statistical lives, does it likewise justify the killing of Sam? It is possible that some acts are justifiable only if, from our ungodlike epistemic vantage point, the risks of an act are borne by many individuals rather than concentrated on one—even if God knows the one on whom the harm will actually fall, and whose risk is therefore one.

⁸⁴ For another illustration, recall Maxwell and Ned (Chart 2). In this comparison, the disparity in risk that each imposes on any single victim is even greater, so it is even clearer that Maxwell's 4,000 drives should not be equated in culpability to Ned's single drive even though the aggregate risk imposed by Maxwell is the same as the risk imposed by Ned. But by the same token, if one can imagine a world in which Maxwell's thousands of drives all endanger only a single victim (a desert island inhabited only by Maxwell and a disabled victim?), then the aggregate risk imposed by Maxwell on that one victim is the same as the risk Ned imposes on his single victim. Of course, Maxwell is still arguably less culpable insofar as the benefits that his 4,000 drives provide far exceed the benefit that Ned's one drive provides.

⁸⁵ A related "concentration of risk" argument also helps explain an otherwise puzzling aspect of self-defense proportionality doctrine.

The defender is permitted to use no more than "proportional" force, but this requirement is weaker than it might seem in several respects. One respect is that the defender is permitted to use individually proportional force against each of multiple attackers, if that is necessary to defend himself, even though the aggregate force he would thereby inflict might seem disproportionate. Consider some examples.

(a) D may kill attackers W1, W2, and W3 who are bent on killing him, even though the loss of their three lives is, in a sense, disproportionate to the loss of his own life.

(b) E may not kill or seriously wound attacker X1 who is bent only on shoving E to the ground, even if this is the only way to stop X1's nondeadly attack. (Suppose E is elderly and frail.)

(c) But F, like D, may use nondeadly force against attackers Y1, Y2, Y3, ... Y10, who are bent only on harming him, even though the aggregate harm he will cause to the ten attackers is, in a sense, disproportionate to the harm that E would otherwise suffer. (Suppose

Another illustration is Eleanor, a character in the introduction. Eleanor is a builder who decides to forge ahead with her tunnel construction project despite her awareness that continuing the project requires lowering a fitting into place and very likely killing a worker trapped below. She would properly be convicted of murder, I believe, for knowingly and unjustifiably causing that death. But why should she be so harshly punished? Would it not be perfectly acceptable for a tunnel builder (remember Agatha) to proceed, with great care, to build a tunnel even if she predicts that this will cause at least one death? And would it not be at worst negligent for a tunnel builder (remember Bertha) to proceed, with less care, to build a tunnel that she predicts will cause at least one death?

The answer is the special stringency principle (SSP): Eleanor has chosen to engage in an act she knows is virtually certain to kill the worker, and that choice requires an extraordinarily compelling justification. The other scenarios are quite different: at no point in time does Agatha or Bertha take a particular action, or decline to employ a precautionary measure, that she knows will very likely kill a particular individual.⁸⁶

that they would have jointly shoved him to the ground; he can prevent this only by use of a stun gun that would throw all of them to the ground.)

(d) Now suppose that G, like F, is fending off a nondeadly attack from many attackers, in this case a hundred: Z1, Z2, Z3, ... Z100. The only weapon at G's disposal will create a 1% risk of causing serious bodily injury or death. Presumably he still may use the weapon; this will still, as to each individual attacker, create only a trivial risk of serious bodily injury. And yet, in aggregate, it is almost certain to cause serious bodily injury or death to at least one of them.

But if G's defensive force is permissible, as I believe it is, then why isn't E's? E, like G, has used the only necessary means to defend himself, and unfortunately has caused serious bodily injury or death.

The answer, I think, is that the proportionality aspect of the right of self-defense does not simply give you the right to use a certain degree of force to defend yourself; rather, it gives you a right to use a certain degree of force vis-à-vis an individual attacker. And you may not "concentrate" a disproportionate degree of force on one attacker, even if you could legitimately "distribute" that aggregate amount of force on a group of attackers, and even if that distributed force would probably result in one of the attackers suffering disproportionate harm.

⁸⁶ The Keating/Scanlon "World Cup" example, noted above, supra note 9=, would be analyzed similarly to Eleanor: even widespread social benefits do not justify failing to rescue the technician, because the broadcaster knows that his omission is highly likely to cause the technician's death.

Frances Kamm offers an example analogous to Eleanor. She supposes that a town has a policy that an ambulance driver will not brake on the way to a hospital if braking will prevent him from saving several lives, even if not braking will kill a pedestrian in his path. In an alternative version, she imagines that we are deciding whether to install a device on an ambulance that somehow will work precisely in the manner that the policy instructs, i.e., *Do not brake if braking will lead to a net loss of lives even if not braking will cause a pedestrian's death*. F.M. Kamm, Intricate Ethics, supra note 19=, at =; F.M. Kamm, Morality, Mortality, Vol. II: Rights, Duties, and Status 303-306 (1996). Deontological principles, she plausibly argues, forbid both the policy and the installation of the device. [= GET QUOTE from intric. ethics =]; F.M. Kamm, Morality, Mortality, at 303 ("[C]ertain ways of causing persons deaths cannot be made permissible by agreement because they are too inconsistent with the moral ideal of the person and interpersonal relations"). A critical reason why the "do not brake" policy or device is impermissible is that the policy/device requires imposing a highly concentrated risk of harm on a single victim.

But this deontological principle is a bit puzzling. Why should it matter that the risk is concentrated in this way? Here are some possible explanations for SSP:

- (a) Highly concentrated risk-exposures are especially likely to be nonreciprocal (disproportionately endangering the victim, and disproportionately benefiting others), and therefore especially likely to be unjust;
- (b) The participants might agree ex ante to accept lower level risks but not highly concentrated risks.

These explanations hardly dissolve the puzzle, however. First, from an omniscient perspective, when workers are accidentally killed in Agatha's tunnel or Bertha's tunnel, isn't the actual risk of death to such a worker 100%?⁸⁷ Second, this approach seems to count some lives as more worthy than others. Why is it more important to save the life of Eleanor's victim than the life of the victim of Agatha or of Bertha? Third, isn't this approach paradoxical? After all, the workers might all prefer to take the chance that they might be the unlucky one who needs to be sacrificed in precisely the manner that Eleanor sacrifices the worker. They might even agree explicitly to take such a chance. And yet many deontologists would still find the sacrifice unjustifiable.⁸⁸ For analogous reasons, deontologists would reject ex ante agreement in other well-recognized scenarios where their principles depart from consequentialist reasoning.⁸⁹

The first and second points can be answered more readily than the third. The question we are addressing is how culpable the actor is in various legal

⁸⁷ See the discussion by Alexander and Ferzan of Sam, at note = supra. See also Richard Brook, *Statistical and Identifiable Deaths*, in *Philosophy and its Public Role* 167 (2004) at 177, asserting that a builder such as Agatha and Bertha "can say ... that, given the evidence, she believes the probability of death to each at risk is minimal. However, assuming determinism, she may also note that whoever dies had no chance to survive." Brook posits other hypotheticals that test the limits of the special stringency principle. For example, suppose someone S knows who will die in the building project but cannot communicate with the builder. If the builder knows that S knows the victim's identity but the builder herself does not know the victim's identity, may the builder continue with the project? *Id.* at 171. Similarly, John Broome asks: "If a definite number of people are going to die, can it really make such a vast difference whether or not it is known who they are?" John Broome, *Ethics Out of Economics* 179 (1999). Broome goes on to offer several examples illustrating how implausible it is to make a project's permissibility depend on whether the victims of the project are known. *Id.*

The short answer to Brook and Broome is that it is much more difficult to justify a highly concentrated risk than a risk that is more widely distributed, and more culpable to knowingly impose the first risk than the second. But I concede that the indirect knowledge (via third parties) examples are difficult to sort out. Compare a case where the builder knows that a manager on the product will knowingly kill one worker (in the manner, say, of Eleanor), where the manager knows the worker's identity but the builder does not. Here, SSP should unquestionably apply not only to the manager but also to the builder if he authorizes the act.

⁸⁸ See Kamm, *Morality, Mortality*, supra note 86=, ch. 11; Leo Katz, *Incommensurable Choices and the Problem of Moral Ignorance*, 146 U. Pa. L. Rev. 1465, 1474-1475 (1998); Scanlon, =.

⁸⁹ Thus, in the famous Transplant Case, it is clearly unjustifiable to forcibly take bodily organs from one person, thus knowingly causing his death, in order to save the lives of five. =cite. Indeed, it is also unjustifiable merely to forcibly take a kidney from A (creating a health risk to him but not killing him) in order to save the life of B. And it would still seem to be unjustifiable to perform such involuntary organ transplants even if all patients in the hospitals agreed in advance to permit this type of utilitarian reallocation of organs and even if the involuntary "donor" pool were limited to such consenting patients. = cite

contexts—for example, how much blame he deserves for purposes of criminal punishment. Even if a deity could foresee which particular victim will die, an actor (such as Agatha and Bertha) who we are thinking of sanctioning lacks that type of foresight when she possesses only statistical knowledge; however, she is indeed highly culpable if she chooses to sacrifice a known victim for the sake of more diffuse social benefits.⁹⁰ And even though the lives of the victims of Agatha and Bertha are as valuable as the life of Eleanor’s victim, the question before us is a different one—viz, what is the culpability of the actor who causes their death? One way to see the difference in perspectives is to ask a different question: if we as a society were in a position to save either two “statistical” victims or one “individually known” victim, we certainly should save the two. Thus, suppose by the same budgetary investment we knew we could either (1) prevent two highway accident deaths (deaths that I assume are caused either by careful drivers such as Larry or by merely negligent drivers such as Maxwell), or instead (2) prevent one red-light runner (such as Clara) from knowingly killing a victim, but we cannot afford to do both. Then we should choose the first investment. Yet it hardly follows that Clara is no more culpable than Larry.⁹¹

The third point, the paradoxical nature of a deontological constraint that disserves even the *ex ante* welfare of the affected persons, is most difficult to explain. Some scholars have tried their hand, in the context either of SSP or of other deontological principles.⁹² Perhaps widespread intuitive responses to examples are as important as arguments in explaining these paradoxical features.⁹³ For purposes of this article, I will simply accept them. Of course, consequentialists will reject this approach. Indeed, if they focus only on the bottom line “body count,”⁹⁴ they will consider Eleanor’s conduct permissible, as well as the conduct of our speeding-to-the-hospital drivers Alfa and Benna.⁹⁵ But most people view this type of unqualified, relentless cost-benefit calculus as morally objectionable, and our legal system appears to reflect that view.

⁹⁰ I have been assuming that we judge whether the risk imposed on the victim is “highly concentrated” from the purely subjective perspective of the actor whose culpability we are evaluating. It is possible to alter this assumption, e.g. to restrict the special stringency principle to situations in which both (a) the actor subjectively believes that he is imposing a highly concentrated risk and (b) a reasonable person (or some other idealized actor) would also believe that the risk is highly concentrated. I do not pursue these complications.

⁹¹ See Katz, *supra* note 85=, at 1474: “[I]f a negligent manufacturer ends up killing fifty-five people in a year and a vicious murderer ends up killing exactly one, and if we could stop only one of them, we would probably try to stop the negligent manufacturer. This would not show, however, that the manufacturer is worse than the murderer.”

⁹² For example, Scanlon offers an individualistic contractarian argument against aggregating risks. For discussion, see Simons, *Tort Negligence*, *supra* note 14=, at =.

⁹³ Kamm famously relies heavily on intuitive responses to hypotheticals in justifying her “intricate ethics.” Kamm, *supra* note 19=.

⁹⁴ I say if they focus only on the bottom line, because consequentialists can include other factors in their social welfare calculus, including the distaste that many citizens would feel if the legal system permitted unadorned “bottom line” calculations of this sort. The standard deontological retort, however, is that when consequentialists revise their cost-benefit calculus by including the cost of disappointing widespread popular preferences for honoring deontological restrictions, they achieve closer alignment of consequentialism with deontology but do so in an unprincipled and unpersuasive way. (For example, this approach achieves closer alignment only to the extent that the policy is known to the populace; so if the policy is deliberately hidden from public view, that might create even greater social welfare; yet it seems illegitimate to defend a policy whose value depends on such secrecy.)

⁹⁵ For example, Kaplow and Shavell strongly imply that they would find Eleanor’s action morally and legally permissible. See note = *supra*.

E. Is the real question whether the victim is identifiable?

This discussion of SSP might suggest that the underlying issue is simply whether the victim is identifiable: if he is, then the actor has individualized knowledge; if he is not, then the actor has statistical knowledge. But this view is oversimplified, indeed mistaken. “Identifiable” has no single meaning in this context, and most of its plausible meanings bear no intrinsic relationship to individualized knowledge. In this section, I hope to show: (1) the strong intuition that “identifiability” is morally and legally highly significant is largely explicable by psychological heuristics and cognitive biases that lack direct moral relevance; (2) however, the intuition is indeed morally and legally relevant insofar as it implicitly reflects the significance of a highly concentrated risk and of SSP.

When the actor has statistical knowledge, the victim is typically unidentified, while if he has individualized knowledge, the victim is in some sense identified. Does it matter to the permissibility or culpability of an act whether the likely or possible victim is identifiable? Intuitively, the answer seems to be yes. Consider the classic illustration: we as a society⁹⁶ are willing to expend far more resources to save an identified, trapped coal miner than we are willing to invest *ex ante* in mine safety in order to save the life of one miner (whom we cannot currently identify) in the future.⁹⁷ Moreover, legal doctrine contains some examples of especially stringent duties to rescue identifiable individuals. First, under traditional tort doctrine, a landowner owes a duty towards adult trespassers only if he “knows” of their presence; his duty is much more restricted if he is aware only of a statistical risk of trespassing.⁹⁸ Second, courts often recognize an unusually stringent duty on the part of a ship’s captain to attempt to rescue someone if he falls overboard, even when the rescue is very unlikely to be successful; indeed, courts sometimes ignore the traditional cause-in-fact requirement in such cases, for they permit full liability even when plaintiff cannot show that but for the breach of duty, the decedent more likely than not would have been saved.⁹⁹

⁹⁶ This implicit policy of differential concern characterizes both government regulatory efforts and private acts of risk-creation and risk-prevention.

⁹⁷ See Calabresi, *supra* note =, *Ideals, Beliefs* at 6; Charles Fried, *The Value of Life*, 82 *Harv. L. Rev.* 1415, 1416 (1969). The classic quotation on the topic is from Josef Stalin: “The death of a single Russian soldier is a tragedy. The death of a millions soldiers is a statistic.” =

In the field of medical care, the identifiability of the victim is relevant in a number of different ways. For example, once a patient is under a doctor or hospital’s care for an existing ailment, a more stringent duty is owed the patient to prevent further harm, more stringent in the sense that greater resources are required to be expended than are required to prevent the patient from developing future illnesses, i.e. to prevent mere “statistical” harms. More generally, “[a] bias toward identified lives can drive priorities in risk management including the choice between prevention and treatment of injury or disease,” Hammitt & Treich, =, at 46. See also Fried, *id.* at 1416; Jenni & Lowenstein, = at 236. See Corso, Hammitt, Graham, Dicker, & Goldie, *Assessing Preferences for Prevention versus Treatment Using Willingness to Pay*, *Med Decis Making* 2002 22: S92 (finding that survey respondents stated a willingness to pay for treatment that was significantly greater than their WTP for prevention, and concluding that “treatment is more strongly preferred by society than prevention when the health context is the same and benefits of each are held constant”).

⁹⁸ =cites, brief discussion.

⁹⁹ See *Gardner v. National Bulk Carriers, Inc.*, 310 F.2d 284 (4th Cir. 1962); =

A quite distinct “identifiable victim” issue arises in tort law: many jurisdictions limit the duty of a therapist to warn about, or take precautions against, a patient’s threats of violence, to “identifiable victims.” =cite. But this limitation is due to the perceived unfairness of imposing broad liability on therapists and the perceived inefficacy of such a

However, the question of the relevance of identifiability is more complex than first appears, because identifiability has various distinct senses.¹⁰⁰ Most of these senses, we will see, have little or no relevance to the issue before us—namely, the culpability of an actor with some awareness (either “individualized” or “statistical”) of the risk he is posing to a potential victim. Moreover, paradigm examples such as the trapped coal miner provide dubious guidance precisely because they combine numerous features, and numerous senses of identifiability, so it is difficult to determine the source of their power over our intuitions.

As a preliminary matter, any analysis of identifiability must be from some perspective—but which one? The most relevant perspective for our purposes is that of the actor whose culpability is being judged, though, in discussing psychological research, I will necessarily have to take the view of ordinary observers, since that is the perspective that the studies explore.¹⁰¹

Consider, then, some possible meanings of identifiability. First, it can be understood as non-anonymity, i.e., a situation in which we have some specific information about the victim. Such identifying information could be very weak—for example, any information that makes the victim other than completely anonymous. Or the identifying information could be stronger—for example, the actor might know the victim’s gender, or age, or name, or might be personally acquainted with the victim.

Second, identifiability could mean salience: the situation of the victim is “on screen,” vivid, emotionally engaging.¹⁰² When the coal miner becomes trapped and in need of rescue, his plight might be televised or otherwise widely reported. But when a large group of miners suffer health and safety consequences due to inadequate advanced health and safety precautions in the mine, they are virtually invisible to public attention.¹⁰³ The Availability heuristic is likely playing a role here: people view a type of event as more probable if it more easily comes to

broad duty in reducing harm. It is not premised on the principles we have discussed, e.g., the supposed greater culpability of one who picks out a named or non-anonymous victim, or the special duty to avoid imposing a concentrated risk on an individual.

¹⁰⁰ For a careful parsing of the difference senses and sources of the identifiable victim effect in the psychological literature, see Karen E. Jenni & George Loewenstein, Explaining the ‘Identifiable Victim Effect,’ 14 J. RISK & UNCERTAINTY 235 (1997). See also Cass Sunstein, Cognition and Cost-Benefit Analysis, in M. Adler & E. Posner, Cost-Benefit Analysis: Legal, Economic, and Philosophical Perspectives 223 (2001) (providing a helpful overview of cognitive biases that affect risk perception).

¹⁰¹ Another relevant perspective is that of an idealized observer; and this perspective is indeed often implicit in the paradigm trapped coal miner example, where the question is often put as what “we” as a society are willing to invest ex post in rescuing v. ex ante in prevention.

¹⁰² Jenni & Loewenstein, =, at 237.

¹⁰³ To be sure, if we examine the two scenarios after a death has occurred, the difference in salience is greatly reduced. If a mine collapses and kills a miner due to inadequate safety precautions, that will provoke enormous public sympathy and scrutiny, perhaps comparable to a failure to save a trapped miner. Nevertheless, even from this perspective, statistical knowledge cases often differ from individual knowledge cases. After the fact of a fatality, it is sometimes difficult to determine, in a statistical knowledge case, whether the inadequate precaution of the actor or instead some background or different risk caused the death, while in an individualized knowledge cases, it is almost always clear that the actor caused the death.

mind.¹⁰⁴ And the emotional connection to a victim is strongest when only a single victim is “on screen.”¹⁰⁵

Third, identifiability could refer to the fact that a victim is already in peril—i.e., the victim needs to be rescued from a peril that has already placed him in a highly vulnerable position. Such a victim can be contrasted with a victim who might be endangered in the future if we fail to take proper precautions. Put differently, when the victim is already in peril, the intervention to save him is in a sense *ex post*; when he is not yet in peril, the precaution needed to save him is (in this sense) *ex ante*.¹⁰⁶ And psychological studies confirm that people react more generously when they confront someone already in peril than someone who is expected to be in peril in the future.¹⁰⁷

Fourth, identifiability could describe an understood, relatively certain causal process for saving the victim. We know we can effectively save the trapped coal miner. We don’t know, with anything like that degree of confidence, that safety precautions would prevent any specified number of deaths or injuries.

Fifth, identifiability could mean that the death is certain, in contrast to a statistically known death, which is uncertain. And, psychologists explain, when we add the fact that people are usually risk-seeking for losses, the result is a tendency to care more about preventing a certain loss (e.g. a certainty that Victor will die) than about preventing an uncertain loss with the same expected value (e.g. a statistical prediction that adopting a particular safety precaution for a large group of miners will reduce the expected number of deaths in the mine from five to four).¹⁰⁸

Sixth, identifiability could be a proxy for proximity in time and space. The trapped coal miner is immediately before us. The miner or miners who would be

¹⁰⁴ See, e.g., Sunstein, *supra* note 97, at 229-230. The evidence in the psychological literature of a distinct vividness effect, separate from the other identifiable victim effects, is surprisingly weak, however. See Jenni & Loewenstein, at 253.

¹⁰⁵ Thus, the study authors found support for their hypothesis that “a single identified victim evokes stronger feelings than an unidentified single victim, or a group of victims, regardless of their being identified or not.” T. Kogut & I. Ritov, The “Identified Victim” Effect: An Identified Group, or Just a Single Individual?, *J. Behav. Dec. Making*, 18: 157–167, 165 (2005).

¹⁰⁶ The *ex post/ ex ante* characterization is a natural one here, but it might be questioned. If the miner in need of rescue has not yet suffered any harm, his plight is also *ex ante*: he is at risk of harm if we do not act now to rescue him. And the same is true of the miner who would in the future be saved from a mine collapse if we do not act now to build more safety features into the mine in which he will work. The key might be that the first miner is “already endangered” or in “immediate, actual peril” (see Fried, *supra* note 97, at 1419) while the second miner is not. But this idea of “actual peril” is difficult to cash out, at least if one views risks as subjective and epistemic.

¹⁰⁷ =cite. Interestingly enough, psychological research reveals that people do respond more generously, in studies examining their degree of altruism, when they know merely that someone in need has been selected as an eligible recipient than when they know that the eligible recipients have not yet been selected. This finding suggests that some type of identifiable victim effect persists even when the person considering a contribution has absolutely no information about the recipients. Deborah A. Small & George Loewenstein, Helping a Victim or Helping the Victim: Altruism and Identifiability, 26 *J. Risk & Uncertainty* 5, 5-14 (2003).

¹⁰⁸ “[T]he number of certain (identifiable) fatalities that is deemed “equivalent” to uncertain (statistical) fatalities is less than the expected number of statistical deaths. Both the certainty effect and risk seeking for losses, therefore, may contribute to the tendency to treat identifiable (and thus certain) victims as more worthy of attention than statistical victims.” Jenni & Loewenstein, =, at 238.

saved by a precaution in the future are not. Just as we feel stronger bonds to those in our geographic community, we might feel a stronger obligation to a person whose plight is immediately before us.¹⁰⁹

Seventh and finally, identifiability could be a proxy for concentration of risk. Here, psychological studies show that participants give much greater priority to a life-saving project if its effects are described as saving a high proportion of a reference group than if they are described as saving a low proportion.¹¹⁰ Perhaps, then, when people show special sympathy towards an identifiable victim, they are responding to an individual victim who is at extremely high risk of death if not rescued.

In the end, identifiability in all of these senses (except for the last) has little to do with the culpability exhibited by individualized knowledge. Moreover, identifiability in most of these senses has at most an ambiguous bearing on the actor's culpability in any other respect. Consider these different senses of identifiability in turn.

First, consider non-anonymity. In the very weak sense, where identifiability merely requires some information that makes the victim non-anonymous, the concept does not even distinguish individual from statistical knowledge. For whenever an actor is analyzing risks that a specified action or inaction would pose, he can readily articulate at least some information about the group that is at risk (future coal miners who might benefit from a mine safety improvement, for example). Even in the stronger sense, requiring significantly more information, the concept is an unhelpful criterion of individualized knowledge. A driver speeding at night-time might perceive only that he is likely to strike a human body; he acts with individualized knowledge that he will harm another person, even if he is unable to determine the age, gender, or size of the person, much less his facial features or name. One could also imagine a mine rescue case in which we were unable to identify the particular person in need of rescue except as one of a large group of miners. Indeed, in a natural disaster, rescuers might come upon a helpless individual whom they cannot identify at all. Yet those who support a stringent duty of rescue are not likely to weaken their support because of this lack of information. And, finally, anonymity of the victim has an ambiguous relationship to culpability: in some respects it makes the actor less culpable, but in other respects more.¹¹¹

Second, greater salience does indeed distinguish the trapped coal miner from the miners who would in the future be saved if we built safer mines. But the question remains whether this differential in the emotional reactions felt by ordinary

¹⁰⁹ Similarly, philosophers have discussed the “nearest and dearest” objection to consequentialism, i.e., the objection that consequentialists cannot account for a perceived greater duty to our neighbors or fellow citizens. = Jackson, *Ethics*; = Kamm, *Intricate Ethics*, supra note 19, Ch. 11 and Ch.12.

¹¹⁰ Thus, if a traffic safety project is described as saving 2 out of 4 persons in the reference class (fatalities at a given intersection) rather than as saving 2 out of 1700 persons in the reference class (traffic fatalities in the state), the first description receives a much higher priority. Jenni & Lowenstein, supra note =, at 252.

¹¹¹ In some cases an actor's willing to harm an anonymous stranger seems to inculcate, for it makes the act more incomprehensible, more frightening, more inhuman. See Heinzerling, *The Rights of Statistical People*, supra note 10, at 196 (noting that in the 1982 Tylenol poisoning incident, it was the very inability to identify who would be victimized that magnified the public's fear; “unidentifiedness is a close cousin of the awful randomness—associated with terrorists and criminal maniacs—that many people uniquely fear”). But in other cases, willingness to harm a person one knows seems to inculcate, for it can express personal spite or malice, or can target a vulnerability in an especially reprehensible way.=

observers should, without more, have normative weight.¹¹² Moreover, salience and vividness are not always present in cases of individualized knowledge, nor are they always absent in cases of statistical knowledge. Consider again the speeding nighttime driver from the prior paragraph, who knows very little about the potential victim; and compare a careful driver who only drives in his own neighborhood, where he knows everyone, including everyone he puts at slight risk, and thus has a vivid and rich appreciation of the lives he is putting in peril.

Third, the *ex post/ex ante* distinction is treacherous to draw in this context, as noted.¹¹³ Moreover, whether someone is in some sense “already in peril” is only indirectly relevant to one’s moral and legal obligations. For example, the burden to save such a person will often, but not always, be less than the burden to take advance precautions to save one life. And again, the *ex post/ex ante* distinction does not precisely track the individualized/ statistical knowledge distinction. For example, if an employer fails to make the workplace safe for her one and only employee, she might have individualized knowledge that the employee will be harmed by her omission, yet, at the time of her negligent omission, the worker might not be in peril.

Fourth and fifth, causal and statistical uncertainty do indeed characterize many cases where advanced precautions might have a beneficial effect on safety levels, and we tend to be less uncertain in these respects about efforts to save or not to harm identified individuals. But these generalizations are crude. Even if the only technique for saving the trapped coal miner was untested, resulting in causal uncertainty, the intuition that there is a strong obligation to rescue him would remain. And although it is sensible to discount probabilistic judgments about the precise effectiveness of future safety measures, it would be folly to ignore strong statistical evidence that has been empirically confirmed.¹¹⁴ Moreover, an actor might have individualized knowledge that he is likely to harm another and at the same time harbor small doubts about the causal effects of his act.

Sixth, temporal and spatial proximity do not really explain the coal miner examples: many Americans would have a similarly strong empathic response to a mine disaster in England than one in another county or state in the United States. And these features do not align very well with individualized knowledge: with modern technology, an actor can knowingly harm individuals from a great spatial or temporal distance.

However, the seventh factor, identifiability in the sense of concentration of risk, is indeed morally and legally relevant, and also might help explain some of the psychological findings noted above.¹¹⁵ In the “trapped coal miner” type of scenario,

¹¹² A policy that gives higher priority to saving more familiar victims “would amount to allowing media coverage to determine aid allocation.” Jenni & Loewenstein, at 240. See also Simons, Dworkin's Two Principles of Dignity: An Unsatisfactory Nonconsequentialist Account of Moral Duties, 90 B.U. L. Rev. 715, (2010).

¹¹³ See note = *supra*.

¹¹⁴ Often people discount very heavily, or simply ignore, low probability risks, a cognitive bias that is classified as Probability Neglect. =cite.

¹¹⁵ Interestingly enough, Jenni & Loewenstein find that of the various factors that might explain the identifiable victim effect, the strongest was the proportion of a reference group that could be saved. = at 254-255. This “reference group” effect is one way of describing the degree of concentration of a risk. Perhaps this psychological tendency is at least partly based on a belief in the deontological principle SSP. On the other hand, the effect is significant in a far wider range of cases than SSP, for it applies even when the proportion that could be saved is much less than 50%. Moreover, the reference group that experimental studies employ in

we do have a concentrated risk, a very high probability that this individual will die if not rescued. Of course, individualized knowledge by definition requires awareness that one is highly likely to cause harm to an individual victim. Moreover, the special stringency principle applies here, at least with respect to those owing a duty of aid to the miner. To be sure, this is an omission case, unlike the previous individualized knowledge and SSP cases we discussed, including Alfa (the actor who knowingly ran over a pedestrian to save the lives of her passengers) or Eleanor (who knowingly crushed the worker in order to continue the tunnel project). However, there is no obvious reason not to apply SSP to omissions by actors owing a duty to the person in need of rescue, as well as to actions.¹¹⁶

But is there anything else to be said in favor of a higher duty to an identifiable victim, other than the “high concentration of risk” argument? Can the other features somehow, in combination, have some significant moral relevance? Perhaps, for example, we must uphold at least symbolically our self-conception as a caring society; and therefore we have a duty to spend much more on persons in need of rescue when their plight is salient and unmistakable.¹¹⁷ On the other hand, perhaps (as I am inclined to believe) this social preference cannot really be defended, but merely reflects an irrational psychological disposition to overvalue salient events, an inability to comprehend the importance of ex ante precautions, and an unjustifiable tendency to discount greatly, or even to ignore, risks of future harm if the probabilities of such harm are low.

However we resolve this debate, I am doubtful that the symbolic argument has great force in the context we are considering, namely, whether an individual actor is more culpable if he knowingly creates a high likelihood of harming identifiable individuals. Even if there is significant expressive value in demonstrating that we are a caring society, that value is only weakly furthered or instantiated by heightening legal sanctions for wrongdoing. To be sure, that expressive value might sometimes be relevant to what actions are permissible in the first place, especially when the actor in question is the government, so that its acts and failures to act have a distinctive expressive dimension. Those attracted to this view might then claim that if a government actor owes a more stringent duty to save identifiable individuals in service of this symbolic value, a private actor also has a

evaluating the effect is susceptible to quite arbitrary framing (id. at 254; and see note = supra) in a way that SSP is not.

¹¹⁶ Many deontologists, to be sure, will recognize weaker duties to affirmatively aid than duties not to harm. But that distinction is irrelevant here. SSP, when applied to omissions, would compare different affirmative duties to aid; it would not compare a duty to aid with a duty not to harm. Thus, SSP would impose a much more stringent duty to aid (a) when the vulnerable person in need of aid is highly likely to die without the actor’s aid (and is highly likely to survive with aid), than (b) when the vulnerable person has only a small chance of dying if not aided.

¹¹⁷ “Decisions that are seen to protect identified victims have the merit of endorsing social values of compassion, solidarity, and unwillingness to sacrifice the few for the benefit of the many.” Hammitt & Treich, at 46. “It may ... be dehumanizing to stand idly by when strenuous, expensive effort has a substantial chance of saving lives.” Allan Gibbard, *Risk and Value*, from *Values at Risk* 94, 101 (Doug MacLean ed. 1986). See also Fried, supra note 97=, at 1425 (describing but questioning the “symbolic value” argument); Ronald Dworkin, *Justice for Hedgehogs* (2010=), at = 178-179, discussed in Simons, Dworkin’s Two Principles of Dignity, supra note 108=, 90 B.U. L. Rev. at 721-724.

To be sure, an alternative, pragmatic argument is plausible: if we fail to recognize a higher duty to the trapped coal miner, the public will lack the political will to invest ex ante in mine safety and other precautions that will save merely statistical victims. But the focus of this article is on the underlying norms of culpability and wrongdoing that a well-informed public should endorse, not on political realities.

higher duty. On this approach, it might be impermissible for a coalmine operator not to expend enormous resources to save a trapped miner, even if it would not be impermissible not to spend equivalent resources in advance on safety precautions to save a statistical life. However, I doubt the premise, that the symbolic value should either increase culpability or make otherwise permissible acts impermissible. In my view, the heightening of duty that SSP demands, and that sometimes requires greater expenditures on saving identifiable lives, is high enough.

F. Why are we tempted to equate statistical with individualized knowledge?

Why, despite the arguments thus far presented, is there a lingering sense that the actor who engages in repeated actions that predictably will lead to a high probability of a harm is about as culpable as an actor who takes a single action that he predicts is likely to cause that harm? Part of the answer is that the arguments we have explored, including ICAA and SSP, are subtle; and the “identifiable” victim effect can seem quite persuasive (even though it often should not be). In this section, I address four other reasons for the ill-advised equation of the statistical and individualized knowledge.¹¹⁸

The first is an understandable confusion between causing harm knowingly (in the individualized sense) and causing harm purposely, and a concomitant assumption that both are equally difficult to justify. But that assumption is false. Causing harm knowingly is much more easily justified. To be sure, knowing infliction of harm is often quite unjustifiable, or justified only by especially compelling reasons such as self-defense and necessity. But this is so only in specific contexts of individual interaction and only with respect to specific types of harm (such as death or serious bodily injury). It is not wrongful knowingly to cause devastating economic harm to another in the course of legitimate market competition, or knowingly to cause severe emotional harm to another as an incident of the break-up of a relationship. In many contexts, knowingly causing harm is both morally and legally permissible. Thus, we need to proceed carefully before assuming that knowingly causing harm should carry a presumption of serious fault

¹¹⁸ A fifth reason, specific to tort law, is a strict liability principle. Under *Vincent v. Lake Erie Transp. Co.*, 124 N.W. 221 (Minn. 1910), an actor can be liable in damages for intentionally destroying the property of another in order to save his own person or property, even if his decision to do so is reasonable (or, in doctrinal terms, even if necessity gives him a privilege to commit the trespass). When an actor engages in a course of conduct that he knows will harm some victims for the greater good of others, or for the actor’s own benefit, then the Vincent principle arguably applies. See James Henderson, Jr., *Coping with the Time Dimension in Products Liability*, 69 Cal. L. Rev. 919, 937 (1981) (discussed in *Armour*, = 1159).

However, once again, this perspective is problematic in not clearly distinguishing an individual act (in this case, a faultless act) that the agent knows will cause harm from a series of acts that the agent “knows” will cause harm only in the statistical sense. If I drive at a reasonable speed and faultlessly harm a pedestrian, I have made an implicit decision to provide benefits to myself (the benefits flowing from speedy transportation) at the expense of risks of harm to others (albeit very low risks). Yet the Vincent principle does not apply to such low-level risks. See [=KWS Culpability and retributive theory?]. =See also Restatement Third draft, comment on strict liability for abnormally dangerous activities (rejecting extension of strict liability to the mere repetition of risky activities).= Still, from a wider perspective, the decision of all motorists to drive rather than walk or bicycle results in a predictable (and probably significant) number of deaths or serious injuries annually even with respect to the class of motorists who drive without any fault.

(or indeed a presumption of any fault), even when the knowledge is individualized (and a fortiori when it is statistical). By contrast, purposely causing harm is virtually always legally impermissible. For example, the Restatement Second of Torts treats the purposeful causation of economic harm as a prima facie tort.¹¹⁹

Those who find abhorrent the law's failure to sanction those who cause harm with statistical knowledge might be unthinkingly confusing statistical knowledge with individualized knowledge, and then individualized knowledge with purpose. This is an understandable confusion in some contexts. We are sometimes epistemically justified in inferring purpose from knowledge (e.g., in concluding that if an actor shoots point blank at the victim and knows that the bullet is very likely to cause his death, he has the purpose to kill). And it is tempting to say that a corporation that is hell-bent on profits, that does not care about the environmental effects of its activities, and that knows those effects will be very harmful, must intend those harmful effects. Indeed, there is empirical evidence that ordinary people often unconsciously employ a moralized conception of "intention": if a known (but motivationally irrelevant) side-effect is harmful, study participants are more likely to characterize it as intended, but if it is beneficial, they are much less likely to do so.¹²⁰

The problem, however, is that knowledge differs from purpose or intention. A person does not always intend the (known) natural consequences of his act. The nurse who gives my child a vaccination does not (I hope!) intend to cause him pain. The woman who regretfully breaks up with her boyfriend does not intend to cause him distress. Even an actor who shoots a bullet that he knows will kill the victim need not intend that result.¹²¹ Recall the introductory examples. Neither Agatha (the careful builder) nor Bertha (the less careful one) intends the injuries that their projects inevitably and predictably will cause, for each has taken precautions against them.

Second, it is also easy to forget that legal standards do not consider an act unjustified merely because the actor is cognitively reckless, i.e. merely because he is aware that he is creating a significant risk of harm.¹²² Modern life is bursting with justifiable risks, and is inconceivable without them. Even the deliberate or purposeful creation of risks is often justified when the risk is not unduly high and the actor creates the risk in pursuit of a socially acceptable benefit (such as the thrill of skiing or the pleasure of observing a trapeze artist walking a tightrope without the protection of a net). Yet cognitive recklessness is the state of mind that most instances of statistical knowledge actually reveal. (Indeed, the bare fact that an actor

¹¹⁹ =cite. Only a few jurisdictions have followed this very broad provision. Simons, *supra* note 30=, at 1084 n. 84. But tort doctrine commonly makes intent to cause certain specified harms tortious, or at least considers an intent to harm a factor militating in favor of liability. =cites.

¹²⁰ See Joshua Knobe, *The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology*, *Philosophical Studies*, 130: 203-231 (2006).

¹²¹ Consider a self-defense case in which the actor's goal is only to completely disable the victim from posing any further threat. If the victim were miraculously to survive while at the same time posing no threat, this actor would be relieved, not disappointed; for he did not intend the victim's death.

¹²² Recall the distinction between cognitive and complex recklessness, *supra* note =; the latter additionally requires that the risk the actor knowingly took was unjustifiable.

is aware of possible harm is what many advocates of statistical knowledge liability most emphasize.¹²³)

Third, the tendency to view statistical knowledge as very similar, if not identical, to individualized knowledge might reflect heuristic error.¹²⁴ Many people have difficulty apprehending the significance of large numbers, and therefore mistakenly assume that if an activity can be predicted to cause even a relatively small number of deaths or injuries, the activity must be excessively dangerous. But that assumption ignores the plain fact that many very widespread or long-continuing activities will inevitably produce harm even if they are conducted with reasonable care. Consider a pair of examples offered by Cass Sunstein:

A. Company A knows that its product will kill ten people. It markets the product to its ten million customers with that knowledge. The cost of eliminating the risk would have been \$100 million.

B. Company B knows that its product creates a one in one million risk of death. Its product is used by ten million people. The cost of eliminating the risk would have been \$100 million.

I have not collected data, but I am willing to predict, with a high degree of confidence, that Company A would be punished more severely than Company B, even though there is no difference between the two.¹²⁵

Sunstein concludes that the response he predicts to these examples illustrate a “moral heuristic,” a rule of thumb (analogous to a cognitive heuristic) that works well most of the time but also sometimes misfires¹²⁶:

I suggest ... that a moral heuristic is at work, one that imposes moral condemnation on those who knowingly engage in acts that will result in human deaths. Of course this heuristic does a great deal of good. The problem is that it is not always unacceptable to cause death knowingly, at least if the deaths are relatively few and an unintended byproduct of generally desirable activity.¹²⁷

Sunstein’s observation is plausible, and it helps explain both the popular sentiment against cost/benefit analysis, and the associated popular inclination to reject ICAA and to judge one who acts with statistical knowledge as harshly as one who has individual knowledge. Of course, the reason that plaintiffs’ trial lawyers would use the rhetoric of the Company A example rather than the Company B example is in order to suggest that (justifiable) popular disapproval of actions taken with individual knowledge should extend to disapproval of actions taken with only

¹²³ See, e.g., LeBel, *supra* note 33=; Henderson, =old article; Heinzerling, *supra* note 10=.

¹²⁴ See Cass Sunstein, *Moral Heuristics and Moral Framing*, 88 *Minn. L. Rev.* 1556, 1578-1580 (2004); Viscusi, *supra* note 12=, at 574 (in one products liability case, a key witness presented evidence of the automobile manufacturer’s conscious cost-benefit analysis; the jury awarded punitive damages, and “they knew” became the “conscious refrain among the jurors interviewed.”).

¹²⁵ Sunstein, *id.* at 1579. I agree with Sunstein’s assessment, though I think the paired comparison would be more fairly presented if the first sentence of the first were phrased less dramatically, i.e. “Company A knows that its product will cause the death of ten people” (rather than “... will kill ten people”).

¹²⁶ *Id.* at 1558-1559.

¹²⁷ *Id.*

statistical knowledge.¹²⁸ At the same time, Sunstein's observation is only a partial explanation of how heuristics affect popular understanding of cases of statistical knowledge, because folk psychology can also undervalue merely statistical (and often temporally distant) risks, at least relative to more salient immediate risks and harms.¹²⁹

And fourth, it is very difficult to develop a criterion for distinguishing the two categories that is both principled and workable. Because of the significance of this problem, it deserves more extensive discussion in the following section.

G. How should we distinguish statistical from individualized knowledge?

Up to this point, the examples we have discussed have straightforwardly fit within either the statistical knowledge or the individualized knowledge category. However, many other examples are difficult to categorize. And sometimes, even when we can confidently classify an example, the implicit criterion that drives our intuitions remains obscure. For example, some cases that intuitively seem very similar to paradigm cases of individualized intention or knowledge (and not merely statistical knowledge) nevertheless involve an extended class of potential victims, or an extended time dimension. An instance is the Time Bomb example, above. Another is the killer who shoots into a crowd, not knowing or caring who he kills, but intending to kill someone. He is just as much an intentional killer as someone who selects a particular victim to kill. A similar pair of examples can be constructed for harms that the actor knowingly rather than purposely causes.¹³⁰ But how do these sorts of cases differ from genuine statistical knowledge cases?

This section evaluates some possible criteria for distinguishing the two categories. It also treats in more detail two issues: (1) an additional "invariant culpability" principle, ICREA, that is sometimes relevant, and (2) a special context, extremely unjustifiable risks, in which the distinction can be ignored.

Consider two recent efforts to develop workable criteria. Neither is quite successful, as we will see, but they both illuminate the problem; and the first identifies a dimension of the problem that we have not yet analyzed.

First, James Henderson and Aaron Twerski investigate the problem in the context of defining "intent" in tort law.¹³¹ Tort doctrine frequently defines "intent"

¹²⁸ One popular strategy of plaintiffs' lawyers seeking punitive damages for product defects is to emphasize the number of injuries that the product caused. Some courts, sensitive to the concerns underlying ICAA, are appropriately careful not to permit punitive damage awards based on number of accidents or injuries alone. See, e.g., *Kopczick v. Hobart Corp.*, 721 N.E. 2d 769, 975 (Ill. App. 1999) (pointing out that the 30 injuries similar to plaintiff's injury from a meat cutting machine represent "roughly 0.5% of the total production of the Model 5700 and 0.0000007% of the estimated total number of cuts made with the Model 5700. Such meager percentages do not put a manufacturer of a mass-produced and inherently dangerous product on notice that its product has an unreasonably dangerous defect."); *Loitz v. Remington Arms Co., Inc.*, 563 N.E.2d 397, 404-407 (Ill. 1990).

¹²⁹ See Wertheimer, comment on restorative justice=; Sunstein, =.

¹³⁰ Thus, suppose that a radical political extremist, with no intention to kill, bombs a particular prominent government building that he knows is occupied at the time by a single victim whose identity he knows. He is a knowing killer. Should we say the same of the terrorist in Time Bomb, who knows only that the bomb will indiscriminately wound or kill random individuals who happen to enter or exit the train station at the wrong time?

¹³¹ James Henderson & Aaron Twerski, *Intent and Recklessness in Tort: The Practical Craft of Restating Law*, 54 Vand. L. Rev. 1133, 1141-1143 (2001).

as encompassing both purpose and knowledge. But how should we define “knowledge”? Henderson and Twerski essentially agree with my claim that knowledge should be confined to what I have called individualized knowledge and must be distinguished from statistical knowledge.¹³² They suggest that the basic distinction is between:

- (1) the harm predicted to result from an individual act;
and either of the following:
- (2a) the harm predicted to result from a course of repetitive conduct; or
- (2b) the harm predicted to result from the interaction of an individual act (one that creates a continuing danger) with the repetitive acts of other persons.

In their view, neither (2a) nor (2b) (two categories of statistical knowledge) should be treated the same as (1) (a type of individualized knowledge).

An example of category (2a) is a baseball player hitting numerous fly balls into the stands over the course of a season. Here, Henderson and Twerski plausibly argue, the player should not be liable (even *prima facie*) for the tort of battery, which requires “intent” to contact, even if he can be practically certain that someone eventually will suffer injury from a batted ball.

Henderson and Twerski’s treatment of category (2a) as instantiating (mere) statistical knowledge, not genuine individualized knowledge, is largely persuasive. This category is usually a straightforward application of the ICAA principle.¹³³ When a baseball player hits a single fly ball, he will at most be aware of only a small risk that the ball will injure a fan (if indeed he is conscious of any such risk at all). And repetition of such acts does not elevate his culpability to the level of individualized knowledge.

But the authors’ category (2b) is especially illuminating: it isolates a new dimension to the statistical knowledge problem. At the same time, for reasons we will now explore, I believe that their specific interpretation of that category is too lenient, treating too many cases as statistical knowledge when some should be treated as individualized knowledge.

¹³² I say “essentially” agree because they frame the question as “whether the consequence complained of must result directly and proximately (both temporal and spatially) from the act.” *Id.* at 1143. In my view, the question is better framed not as a question of proximate cause but as a question of how best to conceptualize “knowledge.” See also Anthony Sebok, Purpose, Belief, and Recklessness: Pruning the *Restatement (Third)*’s Definition of Intent, 54 *Vand. L. Rev.* 1165, 1172 (2001). There is no serious problem of proximate cause when an activity predictably and foreseeably leads to precisely the type of harm that could make the activity negligent in its manner of execution. Criminal law theorist George Fletcher similarly relies on this unpersuasive proximate cause argument to explain why actors such as Agatha and Bertha are not criminally prosecuted. George Fletcher, Rethinking Criminal Law 363 (1978) (arguing that such actors are not prosecuted because of the “implicit rationale” that “the death is simply too remote from the defendant’s act.”).

¹³³ I hedge my endorsement of their category (1), however, because they do not include the qualification I endorse, e.g., that the repetitive conduct must consist of an aggregation of acts that do not differ in any relevant respect from the single act.

1. The principle of ICREA (Invariant culpability when risk-exposures are aggregated)

Here is an example that fits comfortably within Henderson and Twerski's category (2b):

Jungle Gym

A playground equipment company manufactures a jungle gym for use during recess at elementary schools. The company can anticipate that over the fifteen year useful life of the equipment, hundreds of children will suffer minor injuries.

Is the company liable for an intentional tort of battery for causally contributing to these injuries? Not at all likely. The mere fact that repetitive use of the structure by thousands of children is likely to lead to injury does not justify treating the company as negligent or wrongful. It certainly does not justify treating the company as if it had individualized knowledge—as if, for example, it had manufactured an “extreme sport” version of the jungle gym, with slippery surfaces and a 20-foot drop to the ground. In the latter case, the company's knowledge that most child participants are very likely to suffer injury could well result in battery liability.

The reason why “Jungle Gym” should be treated as a case of statistical but not individualized knowledge is quite similar to the rationale underlying ICAA, but it is also subtly different. It is similar because the repetitive use by others of the jungle gym creates, not just a corresponding increase in aggregate risk, but also a corresponding and roughly proportionate increase in aggregate social benefit. Accordingly, whatever judgment we should render about the actor's balance of risk and benefit in one or two instances of use (e.g., permissible, impermissible, grossly impermissible), we should render the same judgment as the instances of use multiply.¹³⁴

However, Jungle Gym is also different from the ICAA cases we have discussed thus far because the repetition of risk exposures and of accidents in Jungle Gym is not due to a repetition of dangerous acts by the company making the equipment; rather, it is due to the fact that, over time, numerous users interact with the potentially dangerous equipment. For completeness, we should recognize a second “invariant culpability” principle. Let us christen it: “Invariant Culpability when Risk-Exposures are Aggregated” (or “ICREA”).¹³⁵ ICAA tells us to presume

¹³⁴ However, if a jungle gym will be heavily used only by one or two actors, arguably the highly concentrated risk would invoke the Special Stringency Principle. Is this a *reductio ad absurdum* of that principle? Is it really more culpable to sell a jungle gym to a family knowing that only two children will use it, than to sell precisely the same equipment to a school knowing that hundreds of children will use it? Perhaps this is the answer: SSP does not apply here because, in both scenarios, the users' consent legitimates the manufacturer's choice to sell the product; consent might be valid here because each risk exposure creates only a very minor risk of injury. On the other hand, consent would not legitimize the product's dangers if the risk on any single occasion of use was very high (e.g. if the user could easily increase the height of the climbable structure to 20 feet or more).

¹³⁵ And we could spell out ICREA in a parallel fashion to the definition of ICAA (with the differences *noted in italics*):

that the culpability of single acts carries over to multiple acts of precisely the same type, while ICREA tells us to presume that the culpability of single instances of risk-exposure carries over to multiple instances of risk-exposure of precisely the same type (even when all of those risk-exposures are the result of a single act).

Now consider a more problematic example of ICREA (or their category (2b)) offered by Henderson and Twerski:

Removed Safety Guard

“[A]n employer permanently removes a safety feature from a machine at a workplace, in order to increase the machine’s productivity... Permanent removal of the safety feature creates a virtual certainty that, sooner or later, a worker will suffer the type of injury that the safety feature would have prevented.”¹³⁶

Henderson and Twerski claim that here, just as in the case of the baseball player, we should not treat the resulting injury as an intended or known result of the act; rather, it should be analyzed under standard negligence principles.¹³⁷

However, the authors’ treatment of Removed Safety Guard as just a straightforward example of their category (2b) or of ICREA, and thus merely an instance of statistical knowledge, is much less persuasive, for two¹³⁸ principal reasons:

Invariant Culpability when Risk-Exposures are Aggregated (“ICREA”):

Suppose: (1) a person engages in a *single act that exposes multiple persons to a risk of harm, such that each risk-exposure is identical* in all respects relevant to legal culpability, including:

the actor’s intentions, beliefs, and motives;

the socially recognized benefits and risks of harm that the act entails; and

the manner in which the benefits, and risks of harm, from the act are distributed among the actor, potential victims, and others;

Then: (2) we should attribute to that actor the same degree of culpability for the whole *set of risk-exposures* as the degree of culpability that we properly attribute to him for any individual *risk-exposure*.

My formulation is broader than Henderson and Twerski’s category (2b) in one significant respect: they seem to require the interaction of the original act with the repetitive acts of others persons, while ICREA only requires that multiple persons be exposed to risks of harm. The broader principle is more sensible, I think. For it encompasses passive victims who do not “act.” (Suppose Fanny builds a bridge that crumbles over time, depositing debris on pedestrians passing underneath.)

¹³⁶ Id. at 1142.

¹³⁷ Id. at 1142-1143. They explain, by way of summary: “for a consequence to be intended in the ‘belief with certainty’ sense, not only must the act producing the consequence be discrete, but the consequence complained of must result directly and proximately (both temporally and spatially) from the act.” Id. at 1143.

¹³⁸ A third possible reason is as follows. Although the initial decision to remove the guard might be justifiable if the machine will remain in that dangerous condition only for a short period of time, the continued failure to repair it might become unjustifiable if lengthier exposure will impose a risk of harm to the workers that exceeds some specified threshold of

First, the safety guard example seems to involve a concentrated distribution of the risk to a relatively small group of workers, many of whom will be repeatedly exposed to the same risk. As we have seen, the ICAA principle loses its persuasive power when the distribution of risk is concentrated in this way rather than randomly and widely dispersed. The same is true of the ICREA principle. Both principles are qualified by SSP, the Special Stringency Principle, which was discussed above. At the extreme, imagine that only one worker operates the machine; if over time the lack of a guard creates a virtual certainty that he will be harmed, his situation is analogous to that of a worker who, the employer can predict, will be harmed the very first time he works with the unguarded machine.

Second, the extended time period also makes it more likely that the owner is aware of the condition and of the opportunity to correct the problem, but chooses to ignore his moral and legal obligation to do so. It might be, in short, a potent illustration of the culpability of defiance, discussed above.

In short, the Henderson/Twerski approach is illuminating, for it reveals that foresight of likely future harm resulting from one's conduct can count as (mere) statistical knowledge even in some instances where the conduct consists of a single act, rather than a series of acts. ICREA is sometimes necessary to supplement ICAA. However, in endorsing an extremely broad version of ICREA, their approach is unpersuasive.

A similar but more persuasive proposed criterion for distinguishing individualized from statistical knowledge appears in the Restatement (Third) of Torts:

The applications of the substantial-certainty test should be limited to situations in which the defendant has knowledge to a substantial certainty that the conduct will bring about harm to a particular victim, or to someone within a small class of potential victims within a localized area. The test loses its persuasiveness when the identity of potential victims become vaguer, and when, in a related way, the time frame involving the actor's conduct expands and the causal sequence connecting conduct and harm becomes more complex.¹³⁹

permissibility. And this is so even if the risks and benefits of a machine lacking a guard both increase in rough proportion. Just as exceeding a threshold level of environmental harm might trigger a more stringent duty of precaution, exceeding a threshold level of risk of personal injury might warrant greater legal protection (switching from no liability to liability, for example, or no longer permitting defenses of consent or assumption of risk).

¹³⁹ Restatement (Third) of Torts, §1, comment e, pp. 8-9.

Moreover, comment e observes, a company's knowledge that its products will, because of the inevitable limits of quality control, contain some manufacturing defects that will cause physical injury might support strict liability, but does not support intentional tort liability. And in other situations, knowledge even of relatively certain harms is consistent with the absence of any tort liability. *Id.* at 9. Thus the comment implicitly endorses ICAA.

The Reporter's Note to this comment shares my criticism of Henderson and Twerski's criterion of knowledge as too narrow:

...Comment e is satisfied if the actor can predict with substantial certainty that someone "within a small class of [persons] within a localized area" will suffer the consequences of the action. Consider the employer that provides employees with a machine that is lacking a crucial safety mechanism; several employees during the workweek use this machine, and the employer knows with substantial certainty that over time one of these employees will suffer injury. Those courts that regard substantial certainty as sufficient to justify the employer's liability for an intentional

This formulation is an improvement over that of Henderson and Twerski, for it offers a straightforward account of the distinction that legal actors (including jurors) can understand and that is broadly consistent with the ICAA, ICREA, and SSP principles.

2. Statistical knowledge of a risk that is extremely unjustifiable

But a significant problem remains. Sometimes “merely” statistical knowledge seems to be as culpable as individualized knowledge. Does this exception swallow the rule? I hope to show that it does not; rather, I believe that the exceptional circumstances in which statistical and individual knowledge are indeed normatively indistinguishable actually “prove” the rule, i.e. they actually underscore the legitimate rationales behind ICAA and ICREA. When the risk is extremely unjustifiable, we will see, it is normally appropriate to treat one who acts with statistical knowledge of the risk as harshly as if he had acted with individual knowledge.

Sometimes we intuitively characterize an actor as “knowingly” causing harm even when it is somewhat indeterminate *ex ante* who will be harmed, or when, or where. Recall this example:

Time bomb, explodes at indeterminate time in near future

A terrorist plants a bomb in a central train station. He knows that the bomb is constructed in such a way that it will explode at an indeterminate time in the near future, indiscriminately wounding or killing many individuals who happen to pass by.

It seems absurd to deny that the terrorist knowingly wounds or kills, and is thus properly sanctioned as a murderer, even though two important dimensions of the act—its timing and the identity of the victim—are not as readily known or knowable *ex ante* as in a paradigm case of knowingly harming another.¹⁴⁰ For it is still the case that when the bomb explodes, it is almost certain to kill whoever is nearby. And the act seems just as unjustifiable and culpable as a bombing in which the terrorist knows which individuals will die.

Now one implicit feature of Time Bomb sets it apart from the types of cases we have been discussing: the terrorist most likely intends to kill or wound. And it is ordinarily proper to treat intended or purposeful harms of this sort as highly culpable, even if the actor believes that the likelihood of success is relatively low. (A terrorist plants a bomb on a plane in order to kill the passengers; he acts with the purpose of killing even if he recognizes that the bomb will probably be discovered.) Purposely causing serious harm is almost always highly unjustifiable. And purpose, unlike knowledge, does not readily partake of degrees; either one has purpose to

tort tend to find that the employer's knowledge in such a case satisfies the substantial-certainty standard.

Reporter's Note, *id.*, p. 16.

¹⁴⁰ See Heinzerling, *Knowing Killing*, *supra* note 10=, at 530 (providing an analytically similar example of an actor who randomly distributes a lethal but undetectable poison in a small percentage of a large population).

cause harm or one does not. Thus, there is no “statistical purpose” problem analogous to the statistical knowledge problem.

But we can easily fine-tune the example so that the terrorist acts with knowledge and not purpose. And once we do so, the example seems similar to more standard individualized knowledge cases. Suppose the bomber’s purpose is not to cause personal injury; rather, his aim is to cause major disruption in the nation’s transportation system. Still, he is undoubtedly aware that he is very likely to cause serious injury or death to one or more people. So modified, the example is not easily distinguishable from a standard individualized knowledge case, such as a driver knowingly running down a pedestrian in order to escape the pursuit of another driver or in order to rush to work. (Recall Clara.) To be sure, the victims of the bombing might be difficult or impossible for the actor to identify in advance. But apart from its questionable legal or moral relevance (as noted above), ex ante anonymity does not distinguish this case from many standard knowledge cases, in which the actor can “identify” the victim in only the weakest sense (for example, the sense in which a driver running a red light can crudely distinguish one complete stranger in the vicinity from another stranger). Nor does an appreciable time lapse between the act of planting the bomb and the explosion causing the harm mark an important distinction. In a standard case, too, serious injuries might occur only well after the actor’s risky act, e.g. a death might occur only well after a car crash. Nor does a significant physical distance between the actor and the instrumentality causing the harm amount to much of a distinction. (Suppose the terrorist detonates the bomb by means of a remote control from another room or even from another city). Again, in a standard case the driver’s act might predictably have some physically distant effects (for example, colliding with a truck transporting chemicals and thus causing widespread environmental harm).

And yet, on closer analysis, even Time Bomb is not a genuine case of individualized knowledge. Individualized knowledge requires some threshold T of probability of harm that the actor believes to exist. Whatever value T has (whether 50%, or 70%, or 95%), in Time Bomb the actor might well believe that the risk to each individual potential victim is much less than T .¹⁴¹

But although this case is not an instance of individualized knowledge, it is most certainly an instance of very high culpability. The actor knows that someone will die as an immediate result of his act. And his act is grossly unjustifiable. So it is not surprising that we intuitively wish to treat him as if he acted with individualized knowledge.

Now recall the Removed Safety Guard example, where the employer removed a guard from a machine and knew that one of a small number of workers would almost certainly suffer injury as a result. This case, too, is probably not a genuine instance of individualized knowledge. Suppose the employer has five workers each of whom uses the machine about the same amount of time, and the employer can predict that there is only about a 20% probability that any one of them will suffer harm, but a high likelihood that at least one of them will. The justifiability of his conduct is not much different than if he removed the guard while employing only one worker, knowing that the worker was almost certain to suffer

¹⁴¹ For another example, compare: (1) an actor who plays involuntary Russian Roulette with a single victim, firing multiple shots which in aggregate impose a risk greater than T , and (2) an actor who plays Russian Roulette with multiple victims, firing multiple shots which in aggregate impose a risk greater than T , but firing only one shot at each victim, each imposing a risk less than T . The first shooter acts with individual knowledge; the second does not. See the Russian Roulette examples in the Appendix.

injury. Again, it is plausible that we might wish to treat Removed Safety Guard the same as a case of individualized knowledge for purposes of legal sanctions—and that is what the Restatement Third of Torts provides, as we have seen.

In short, Time Bomb and Removed Safety Guard show that it is entirely defensible for the law to endorse a relaxed, capacious definition of individualized knowledge, applying it not only to paradigm cases involving a probability of greater than T that an individual victim will be harmed, but also to cases involving small numbers of individuals when the actor is certain (or is aware of a probability greater than T) that one of them will be harmed. On this approach, there is no bright-line distinction between individualized and statistical knowledge. And for similar reasons, SSP is probably best understood, at least for purposes of legal sanctions, as imposing (1) a duty that becomes continuously more stringent as the concentrated risk that the actor believes he is creating increases, rather than (2) a special duty that is triggered only when a defined threshold risk of harm is reached.¹⁴²

So far I have examined cases that stretch the scope of individualized knowledge only modestly. But in other cases, the indeterminacy of time, space, and possible victims is much greater. And yet, in some of these cases, too, cases that are more aptly described as involving statistical knowledge, we have reason to treat the actor as harshly as we would if he had individual knowledge. Consider an example offered by Jody Armour, which, he suggests, shows that a criterion of intention (and, by implication, a criterion of individualized knowledge¹⁴³) can legitimately apply even over an extended period of time:

¹⁴² Nevertheless, even if SSP is a continuous function, it is not a simple proportionate function. That is, as the risk to a single victim, or to a group of potential victims, becomes more concentrated, other features of the actor's conduct that are relevant to its justification must become more compelling, not just in a proportionate way, but in a more dramatically escalating way. For an illustration, recall Alfa and Benna. Suppose it is permissible for Benna to impose a 20% risk of death in order to save one life, but impermissible for Alfa to impose a 100% risk of death even in order to save five lives. Now imagine two intermediate cases. Charles imposes a 60% risk of death in order to save three lives. Dennis imposes an 80% risk of death in order to save four lives. On an invariant threshold version of SSP, it is triggered only when the risk exceeds a particular level, such as 75%; on that version, Dennis acts impermissibly, while Charles, like Benna, acts permissibly. But on a more plausible continuous version of SSP, Charles (as well as Dennis) might act impermissibly even though Benna acts permissibly. (Indeed, when we consider the different degrees and types of harms that risky conduct can cause, it is quite doubtful that any invariant threshold version applies in all contexts.)

One objection to SSP is that it is unnecessary. Instead, we should simply apply a sliding scale, increasing the requisite scrutiny or the requisite weight of justifying factors when the probability of the harm (as perceived by the actor) increases. That is the approach endorsed by Alexander and Ferzan, *supra* note =, at =. Although superficially attractive, this strategy is not an improvement. First, simply indicating that a range of relevant factors are relevant and requiring them to be balanced in a sliding scale provides an opaque formula. See Simons, *Retributivism Refined - Or Run Amok?*, 77 *U.Chi.L.Rev.* 551, = (2010). And, as articulated by Alexander and Ferzan, the formula takes the form of a very simple function of risks and benefits, a function that cannot explain the differing evaluations of Alfa and Benna. An adequate explanation requires a more complex function such as SSP.

¹⁴³ Armour's original example apparently involves "intention" to injure in the sense of purpose, for he assumes that the shooter deliberately aims the gun at a crowd. But I have modified the example to encompass an actor who lacks that intention but possesses individualized knowledge that he will cause injury. The entrance to a train station, the shooter knows, is likely to be extremely crowded most of the day. But the shooter need not intend to harm or kill anyone; his purpose might simply be to cause public panic.

Programmed Rifle¹⁴⁴

A shooter aims a rifle at the public entrance to the train station. The rifle contains a hundred randomly selected rounds, only one of which is “live.” The rifle is connected to an automatic timer and is programmed to shoot once per day or week.

Notwithstanding the significant time gap between the discrete acts, Armour claims, “our intuition would demand that he be responsible for an intentional injury when the ‘live’ round is finally discharged into the crowd.”¹⁴⁵

Armour’s example is indeed persuasive, and it demonstrates the need to add a qualification to the accounts we have discussed, including that offered by Henderson and Twerski. Two features of his example seem important—(1) the actor’s precommitment to a course of activity that will cause serious harm, and (2) the extremely unjustifiable nature of the conduct. In the end, I think the second feature is the critical one in formulating a defensible criterion of individualized rather than statistical knowledge.

Consider first the actor’s precommitment to a course of action. Precommitment of different types and degrees can be distinguished,¹⁴⁶ but let us begin with the clearest case: where the actor makes a single decision that (without further effort by the actor) will reliably result in future imposition of risk. “Programmed Rifle” is such a case, but so is the decision of a corporate executive to market a product (e.g., a knife or an automobile) that will undoubtedly causally contribute to injuries to many of its users. On reflection, this cannot be a decisive factor. Businesses often, and justifiably, precommit to the marketing of a large number of products over an extended period of time, knowing that some of the products will inevitably cause harm. The builder of the jungle gym knows that innumerable children will use it over the course of many years, and some will suffer injuries. Individuals also implicitly precommit in a comparable way, by deciding to engage regularly in a risky activity, such as driving a car or skiing. To be sure, a single individual might not be statistically certain that, by participating in such activities, he will cause serious harm, even over his lifetime. Still, a group of such individuals could be statistically certain of such a result. Yet it would hardly be proper to treat, say, the American Automobile Association as an “intentional” or knowing wrongdoer because it encourages its members to drive—to drive responsibly, but indeed to drive.

¹⁴⁴ The example is adapted from an example in Jody Armour, *Interpretive Construction, Systemic Consistency, and Criterial Norms in Tort Law*, 54 *Vand. L. Rev.* 1157, 1160-1161 (2001). Armour offers it in criticism of the criterion, just discussed, proposed by Henderson and Twerski.

¹⁴⁵ *Id.* at 1161.

¹⁴⁶ In the strongest form of precommitment, once he has performed the initiating act, it is impossible for the actor to prevent the harm from occurring. In a slightly weaker form, it is possible but extraordinarily difficult to prevent the harm (i.e., it would not be reasonable to expect the actor to do so). In a yet weaker form, the actor could and should prevent the risk from taking effect, as when a product manufacturer should recall the defective product. (When an automobile manufacturer sells a defective vehicle, it could exhibit any or all of these forms of precommitment.) And customary patterns of conduct could also be described as a weak form of “precommitment”: I habitually drive to work rather than walk or bicycle or take a bus. But these distinctions are not critical to the main argument in the text.

This is not to say that precommitment is never relevant to legal liability. If Domino's Pizza had a policy (as some alleged) to encourage its delivery drivers to speed if necessary in order to deliver pizzas to customers within 30 minutes,¹⁴⁷ then their knowledge that this policy would lead to a predictable number of accidents and injuries is highly relevant to their liability. And if that policy was sufficiently unjustifiable, and if Domino's knew of its effects, punitive damages would be appropriate.¹⁴⁸ But that conclusion would be true whether drivers followed the policy once or a thousand times (assuming no difference in the company's knowledge of the risks or in other relevant factors).¹⁴⁹ In short, the ICAA principle seems to apply: repetition of conduct that conforms to the policy does not, by itself, make the conduct more culpable.

Now consider the second feature in the "Programmed Rifle" example—the clearly unjustifiable nature of the risk. This feature does seem to be a compelling reason for classifying such cases within the individualized knowledge rather than statistical knowledge category. The fact that the actor maliciously intends serious injury, or at the very least knows that such injury will occur and is moved to act by a sadistic or other highly immoral reason, very clearly distinguishes him as among the most culpable wrongdoers. This is true even if no one can know, when he precommits, the precise time or the precise victim of the future injury. Indeed, one could extend the spatial and temporal dimensions of the example indefinitely, yet we still would want to treat the actor as harshly as an intentional or (individually) knowing wrongdoer. Consider this case:

Atmospheric Poison

A terrorist sends into the atmosphere a specially-designed chemical that he knows is virtually certain to randomly kill a hundred people somewhere on Earth some time in the next hundred years.

Even if the terrorist does not intend the deaths—suppose, for example, that he merely intends to frighten those currently alive, and does not care whether or not the poison has its expected future effect—his clear lack of justification and immoral motivation warrant treating him as extraordinarily culpable, just as culpable as one who poisons a local water supply expecting 100 people to die as a consequence within the hour.¹⁵⁰

¹⁴⁷ Cf. *Wauchop v. Domino's Pizza*, 832 F.Supp. 1572 (D. Ind. 1993) (policy of Domino's Pizza guaranteeing pizza delivery within 30 minutes allegedly contributed to death when delivery driver struck victim's minivan; however, policy explicitly required drivers to comply with traffic laws).

¹⁴⁸ Moreover, if the repetition of the conduct of speeding was pursuant to official policy, rather than simply a pattern of conduct of which the company was aware, that could, of course, bear on the company's legal responsibility under respondeat superior.

¹⁴⁹ In practice, of course, when an actor repeats his conduct, he will often have a more confident and certain belief about its harmful results; and similarly, a legal decision-maker will often be more confident in her judgment that the actor was aware of the high risks he was posing. See TAN = supra. Moreover, it is especially culpable to persist in posing risks that one knows or should know are unjustifiable; this is an aspect of the concern about "defiance," noted above.

¹⁵⁰ Note that this example involves a single culpable act that exposes numerous people to risk. Thus, it exemplifies ICREA, not ICAA.

What is the deeper explanation for this “extremely unjustified” qualification? When we recall the basic rationale for presumptively treating agents as serious wrongdoers when they act with individualized knowledge that they will cause harm, the reason for the qualification becomes more evident. In a standard individualized knowledge case—for example, when a driver realizes that in running a red light he is very likely to run over a pedestrian—the actor is especially culpable because he chooses to act despite his knowledge that he is creating a very high risk of a direct, immediate harm to a victim. Ordinarily that decision will be extremely difficult to justify, given his beliefs and intentions as well as the reasons that should have guided him. Even if the driver’s reason for so acting is legitimate, he is not justified if that reason is not very weighty (if, for example, his speeding through the light merely serves his convenience, in getting home earlier). Compare an analogous statistical knowledge case. Now, an actor is justified in driving (at a reasonable speed) rather than walking home, even though he realizes that this creates a very small risk of serious physical injury to random pedestrians, bicyclists, and occupants of other vehicles, because those same benefits are now sufficiently weighty and the risks of the activity are low and are dispersed among numerous potential victims.

But in both the Programmed Rifle example and the Atmospheric Poison example, there are no legitimate compensating benefits whatsoever. The analysis is dramatically different if we modify the example to one in which the chemical will produce countervailing benefits:

Atmospheric Cancer Cure

Scientists release into the atmosphere a chemical which they believe will have highly beneficial effects in reducing the incidence of cancer from air pollution, but which at the same time will, as an unavoidable side-effect, have the same expected mortality effect as the terrorist’s poison: it is virtually certain to randomly kill a hundred people somewhere on Earth some time in the next hundred years.

Even if the scientists overestimate the compensating benefits, and even if their error is negligent, they would not and should not be treated as harshly as intentional (i.e., either purposeful or knowing) wrongdoers, such as the terrorist releasing the atmospheric poison.¹⁵¹

Another way to appreciate the critical importance of this justification issue is by examining the significant doctrinal distinction between (individualized) knowledge and recklessness. Knowingly causing harm is much more difficult to justify than recklessly doing so. Suppose that “knowingly” causing a particular harm (such as serious bodily injury) is somewhat arbitrarily defined as requiring a subjective belief that one’s act creates at least a 70% risk of that type of harm. Then an actor causes that harm “recklessly” rather than “knowingly” if she subjectively

A separate question is whether future deaths should be discounted to present value, so that killing 100 over an extended future time period is no worse than killing, say, 70 today. However this question is resolved, the main argument here is unaffected.

¹⁵¹ But suppose the terrorist honestly believes that the social benefits from his action (obtaining justice for his ethnic or religious group, destabilizing an oppressive regime, etc.) are as great as those achieved by the atmospheric cancer cure? The short answer is that legal and moral norms trump the individual’s subjective beliefs about the legality and morality of his conduct. I cannot explore the complexities of this issue here.

believes that the risk of harm is less than 70%.¹⁵² There will, of course, be close cases. But it is sensible for the law to employ a relatively bright-line rule that only a belief that one is creating a high degree of risk will trigger the distinctive doctrinal consequences of knowing or “intentional” harms identified above (especially (1) greater punishment or liability, and (2) the obligation to demonstrate a weightier justification in order to avoid all liability). Murder (knowingly causing death) is justifiably punished more harshly than reckless manslaughter; “intentional” battery (including knowingly causing harm) is justifiably treated as a more serious tort, in several respects, than negligence (causing harm recklessly or negligently). Although borderline differences in the perceived risk of harm (e.g., 75% v. 65%) do not so clearly warrant differential consequences of this sort, more substantial differences (e.g. 75% v. 10%) do.

And why are these differential consequences warranted? First, as noted above, ordinarily the greater the risk of harm knowingly created, the more the actor has impermissibly devalued the interests of the victim (or other socially valued interests) relative to the actor’s own ends. Second, we normally cannot be sure that the actor who knowingly creates a smaller risk in pursuit of his ends would have been willing to knowingly create the greater one in pursuit of the same ends. And third, all else being equal, one who knowingly creates a higher risk needs a weightier justification in order to deserve exemption from legal responsibility.

All three reasons argue in favor of treating knowing actors more harshly than reckless ones. But none of these reasons argues in favor of treating actors with statistical knowledge as harshly as those with individual knowledge, as early discussion confirms. (Consider the third reason again: statistical cases simply involve aggregation of numerous low-risk individual cases; and if a low burden of justification is appropriate for the latter, it should be appropriate for the former.) And yet, all three reasons do militate in favor of treating a small subset of statistical knowledge cases, such as Time Bomb, Programmed Rifle and Atmospheric Poison, in the same way that we treat individualized knowledge cases. For if there is absolutely no socially acceptable justification for creating the risk, then:

¹⁵² In the text, I assume that “reckless” is understood in the purely cognitive sense, see note 28= supra, so that we are indeed comparing apples with apples. Most criminal law and tort “recklessness” criteria are more complex than this: they require, not only a cognitive awareness of a risk of harm, but also the taking of an unjustified, or clearly unjustified, risk. See *id.* This common use of a complex rather than purely cognitive conception of recklessness accentuates the legal difference between knowledge and recklessness: if D1 knowingly causes a serious harm, D1 is typically liable unless he can invoke an affirmative defense (such as self-defense or necessity); but if D2 recklessly causes a serious harm, D2 is typically liable for recklessness only if the state (or the plaintiff) can also show that D2’s act was unjustified (or clearly unjustified). For affirmative defenses, a criminal defendant must at least produce supporting evidence before the burden of proof shifts to the prosecution; and sometimes, the defendant has the burden of proof to disprove the defense. See Dressler, *Understanding Criminal Law* 77 (5th ed. 2009); MPC §1.12(2) & MPC §1.12(2). Tort law typically shifts the burden of persuasion to defendants for affirmative defenses. Dan B. Dobbs, *The Law of Torts* 36-38 (2000).

A jurisdiction also might narrow the definition of recklessness by requiring a subjective belief that the risk is above some particular threshold, e.g. 1% or “substantial.” It seems that the Model Penal Code imposes a requirement that the actor believe that the risk is substantial, though the matter is not entirely clear. See Simons, *Should the Model Penal Code be Amended?*, supra note 25=, at =. But whether or not a jurisdiction narrows recklessness in this way, the basic analysis in the text comparing knowledge and recklessness remains valid.

- (1) The actor has very significantly devalued the interests of the victim and of society relative to the (nonexistent) value of any permissible interest the action serves;
- (2) The actor often has revealed that he was willing to create an even higher risk if that had been feasible,¹⁵³ and
- (3) The actor has made irrelevant the question whether he has imposed a sufficiently high level of risk to warrant a strong presumption of unjustifiability, for by hypothesis, his conduct is clearly unjustifiable.

Put differently, a principal reason why the law isolates individualized knowledge that one will cause harm as an especially culpable state of mind is that acting with such knowledge is normally not simply unjustifiable, but especially unjustifiable. This is not usually true of acting with statistical knowledge. Yet in a subcategory of the latter cases, we have independent reason to conclude that the conduct is especially unjustifiable—for example, because the conduct is prompted by a highly immoral motivation).

Criminal law doctrine does recognize this subcategory of cases in its “depraved heart” or “extreme indifference” murder category. In such cases, an actor who would otherwise be considered only reckless is deemed to be as culpable as one who knowingly or purposely causes a death.¹⁵⁴ One who plays Russian Roulette with the victim is guilty of murder if he spins the cylinder and happens to discharge a live bullet on the first pull of the trigger, even if the probability of death is only one in six.¹⁵⁵ Such an actor clearly lacks individualized knowledge, but his depravity is sufficient to justify an extremely severe sanction. The actors in Programmed Rifle and Atmospheric Poison could readily be convicted of murder under this same doctrine.

To be sure, the “depraved heart” murder doctrine itself is controversial insofar as it permits very severe punishment premised on very uncertain criteria. But, for our purposes, the doctrine does reveal that legal doctrine can and does recognize certain “exceptions” to usual mental state criteria when those exceptions “prove,” or serve the underlying rationales of, the usual rules.

Accordingly, the Restatement Third’s criterion for distinguishing individualized from statistical knowledge is essentially sound: individualized knowledge should indeed normally be restricted to “knowledge to a substantial certainty that the conduct will bring about harm to a particular victim, or to someone within a small class of potential victims within a localized area.” But the criterion must be qualified to account for unusual cases such as Programmed Rifle or Atmospheric Poison, with a proviso such as the following: “or, in rare cases, an actor’s knowledge that he will harm someone within a more diffuse or temporally remote group of victims, in circumstances where causing that harm is extremely

¹⁵³ But not always. Consider the sadist who wants to badly injure but not kill the victim; his dangerous acts do not demonstrate a willingness to kill, but he is still about as culpable as someone who is indeed willing to kill. See Kenneth W. Simons, Does Punishment for “Culpable Indifference” Simply Punish for “Bad Character”? Examining the Requisite Connection between Mens Rea and Actus Reus, 6 *Buff. Crim. L. Rev.* 219, 266 (2002).

¹⁵⁴ The Model Penal Code explicitly requires both recklessness as to the risk of death and “extreme indifference.” §210.2(1)(b). Other jurisdictions are less clear about whether recklessness as to the risk must always be proven.

¹⁵⁵ Even if the defendant credibly commits to pull the trigger only once, many jurisdictions would uphold a conviction of extreme indifference murder.

unjustifiable.”¹⁵⁶ Such a proviso permits the law to treat the culpability displayed in these unusual cases as comparable to the culpability displayed in paradigm individualized knowledge cases.

IV. Applications

The analysis above is illuminating in a range of legal contexts. This section discusses three.

A. Knowledge in contexts other than knowingly causing serious harm

“Knowledge” that your action is very likely to cause harm is not always culpable. It is only presumptively culpable when you have individualized knowledge. Subject to rare exceptions, only in this situation is the knowing actor very likely to be acting in an unjustifiable way. How widely do these conclusions apply? Do they apply beyond the context we have been discussing, actors who knowingly cause death or serious bodily injury? Do they apply when the knowledge in question pertains to a legally relevant circumstance, rather than to a legally required causal result?

¹⁵⁶ This solution does not resolve all conceptual difficulties. One might object that it proves too much: it seems to imply that creating even a very low-level risk (one that would normally be categorized as reckless or negligent at worst) can count as individualized knowledge, and should accordingly be sanctioned especially harshly, if the actor’s reason for creating the risk is sufficiently immoral or unjustifiable. Thus, suppose you drive carefully, and you take a drive not for pleasure, convenience, or to accomplish some socially acceptable goal, but merely because you sadistically hope that your careful driving will nevertheless injure or kill someone. See Alexander & Ferzan, *supra* note =, at = (discussing a similar example, Frankie). My analysis seems to imply that if your improbable plan is realized, you should be treated as having knowingly harmed or killed the victim. Arguably the driver is on a moral par with someone who plays Russian Roulette with an unwilling victim employing a device that has a 1 in 1,000,000 chance of killing him.

It is difficult to have confidence in one’s intuitive response to such strange examples, but here are some reactions. An attempt to kill another with a gun should clearly be sanctioned harshly (whether or not it succeeds) even if the subjective and objective risks of success are only one in a thousand. But it is much less clear that the careful but sadistic driver should be treated harshly, or indeed sanctioned at all. The example raises evidentiary doubts: if he were really such a sadist, obtaining pleasure from endangering others, why didn’t he just deliberately run someone over, or at least deliberately drive dangerously rather than carefully? The causal nexus here is also likely to be less direct or foreseeable: careful driving will not cause injury unless unusual events or mechanisms intervene. And the fact that the actor’s conduct conforms to standards of reasonable care makes it much more difficult to credit proof that his actual motivation was an illicit reason. Moreover, if the risks he poses truly are extraordinarily minimal, it is possible that those exposed to them implicitly consent to them, or that there is some very modest social benefit to his creating the risk. See Simons, *Retributivism Refined*, *supra* note 139=, at 565 (analyzing a similar character, Frankie).

On the other hand, if that proof were strong enough, perhaps we could indeed justify harsh sanctions for the careful driver who kills. Thus, suppose Leo wants to kill Larry, but he knows that he, Leo, is under constant police surveillance. Yet he also knows that Larry often sleepwalks in the completely dark street outside his house at 2 a.m., oblivious to his surroundings. If Leo repeatedly drives by Larry’s house at 2 a.m. at a very cautious rate of speed, unable to see whether anyone is on the street, but hoping thereby to kill Larry, and if on the 500th drive he does cause Larry’s death, it seems appropriate to treat Leo as a murderer. And if there really is equally strong proof in other cases that the actor took a (careful) drive simply because he wanted to endanger or kill another, perhaps he, too, should be treated harshly.

We must be cautious before applying the analysis to other “knowledge” requirements in the law. Consider the following legal prohibitions that require knowledge before liability will attach:

1. A criminal defendant must (to be liable for murder) know that his actions will cause the victim’s death.
2. A tort defendant must (to be liable for battery) know that his action will cause offense or physical harm (or must know that his action will cause a nonconsensual physical contact).¹⁵⁷
3. A business owner must know that he is hiring an undocumented immigrant, or a person under the legal age; or that he is selling alcohol to a person under the legal age.
4. A business or individual must know that the property that it received from another is stolen.
5. A lawyer must know that his client will give false testimony.

In each case, the law’s prohibition is very likely intended to encompass individualized knowledge but not mere statistical knowledge. Presumably the laws against knowingly putting on a witness who will commit perjury focus on whether the lawyer is aware of a very high probability, in a single trial, that a witness will lie.¹⁵⁸ Laws prohibiting receipt of stolen property focus on an actor’s awareness that a particular item of property is probably stolen, not his awareness that some item or other in his possession might be stolen. In short, some version of the invariant culpability principles most likely applies in all of these cases. An actor’s legal culpability normally should not change just because his business is high volume or his activity extensive.

However, the five examples may well differ in just where they do and should draw the individualized/statistical knowledge line. In this respect, the first two contexts differ significantly from the others. In these two contexts, the wrong is serious and of a fundamental sort. A moral wrong underlies the legal wrong,¹⁵⁹ and partly for this reason, SSP is an important part of the explanation for presumptively treating action with individual knowledge as especially difficult to justify.¹⁶⁰

But in the other contexts, a serious moral wrong need not underlie the legal wrong. Moreover, the specified legal wrong might simply be a proxy, in useful rule-

¹⁵⁷ Actually, although the requisite “intent” for battery encompasses knowledge, courts differ concerning the requisite object of that intent. Some require merely an “intent” to contact another person; others also require an intent to harm or offend. See Simons, *A Restatement Third of Intentional Torts?*, supra note 30=, at 1066-1070.

¹⁵⁸ Research has not disclosed any judicial definitions of “knowledge” in this context that refer to probabilities.=

¹⁵⁹ In criminal law terminology, it is *malum in se*, not *malum prohibitum*.

¹⁶⁰ However, the second example, the tort of battery, encompasses some relatively minor harms as well as more serious one, for it embraces unpermitted contacts that merely cause offense, not physical injury. SSP might not apply to these less serious instances of battery, just as it might not apply to examples three, four, and five.

Moreover, one of the reasons for clearly distinguishing statistical from individualized knowledge relates to the grading role of the latter relative to other available grades of an offense, such as recklessness or negligence. But in many legal contexts where a knowledge requirement is employed, no lesser grade or degree of culpability is recognized. (This is likely true of examples three, four and five listed in the text.)

like form, for the real legal or moral wrong that the legal standard is meant to address. Consider the third example, a law against selling alcohol to a minor. This criminal law rule is a pragmatic proxy for an underlying norm such as, “Don’t endanger the health of immature people who do not fully appreciate that they are jeopardizing their own health and safety.”

Why do these distinctions between the first two categories and the other categories matter? For three reasons. First, some of the arguments that we have explored, especially SSP, do not apply, or do not apply with the same force, to the last three examples. The concern about not imposing highly concentrated risks on a single victim is much more persuasive when the risk in question is serious personal injury or death, than when the risk is that a trial will be distorted by false testimony or that a person’s property will come into the possession of a person with no right to it.¹⁶¹ Stated more generally, when the actor’s knowledge pertains to a legally relevant circumstance rather than to a legally required result, it is quite unlikely that SSP applies.¹⁶²

Consider:

Pawn shop owners P1, P2, and P3

Owner P1 receives and offers for sale a particular piece of jewelry that he knows is stolen. P2 adopts a sloppy policy for screening goods received by his pawn shop that he knows will result in 10% of the retained goods being stolen. P3 adopts a more effective screening policy that he knows will result in only 1% of the goods being stolen. P2 and P3 each receives and offers for sale a particular piece of jewelry, not realizing it is stolen.

P1 clearly acts with individualized knowledge, while P2 and P3 arguably do not. However, these distinctions are not readily justified by SSP, with its special concern that an actor not impose highly concentrated risks on individual victims. To be sure, in all three cases, the actual jewelry owner is directly harmed when a pawn shop improperly takes possession of his property. But it is doubtful that an especially stringent deontological constraint applies when one creates, say, a 70% risk of improperly depriving a victim of her property yet does not apply when one creates a 30% or 10% or 1% risk.

Second, in cases such as the third example, where a rule serves as a proxy, the judgment about whether individualized knowledge or some lesser culpability should suffice is more complex than in the “knowledge of result” cases canvassed earlier in this article: sometimes a lesser culpability is perfectly consistent with the underlying norm.¹⁶³ Third, in a particular legal context, we might have good reason to employ an especially broad or narrow definition. If, for example, it is important to ensure that the insurance exclusion from worker’s compensation for “intentional”

¹⁶¹ For a series of illustrations, see Appendix, at = .

¹⁶² For a possible counterexample, however, see the example of actors S1, S2, and S3, in the Appendix.

¹⁶³ For example, requiring only reckless awareness of a risk that the person purchasing alcohol is under age might be consistent with properly respecting the underlying norm not to contribute to the health problems of an immature person.

torts remains narrow (or broad), then we have ample reason to define the individualized knowledge prong of “intentional” torts narrowly (or broadly).¹⁶⁴

Still, for practical reasons, it is not surprising that legal standards in quite different domains (criminal law, tort, regulation) and in different subdomains (criminal homicide v. criminal theft) often employ similar mental state categories. The proper definition of knowledge is complex enough as it is; legislators, judges, and other legal actors have reason to use as uniform a criterion as is practicable in different contexts.

- B. If a cost-benefit analysis supplies knowledge that an activity will cause harm, is it culpable to proceed?

The propriety and legitimate scope of cost-benefit analysis, whether conducted by government regulators or by private actors, is an enormous topic.¹⁶⁵ For purposes of this paper, I focus on one question within the host of issues that the topic raises. If an actor conducts a serious cost-benefit analysis of whether to take a precaution or alter his activity, he will come to appreciate many of the harmful consequences of his action or course of conduct. If he then proceeds with the action, is he presumptively culpable, or even highly culpable?

The short answer is no. The longer answer comes in six parts.

First, to presume culpability in this context is unwarranted, for familiar reasons. The invariant culpability principles (ICAA and ICREA) tell us that expanding the scope or duration of an activity does not, for that reason alone, make one more culpable. Furthermore, if we did treat predictions of harm derived from a thorough cost-benefit analysis as individualized knowledge, the result would be dramatic and unacceptable; for it would be extremely difficult if not impossible to justify such knowing killings under existing legal criteria for the necessity or lesser evils defense.¹⁶⁶

Second, in this context, it is understandable that people mischaracterize statistical knowledge cases as individualized knowledge cases, for reasons we have explored. Thus, when the causal effect of which the observer is cognizant is a serious harm, people often (but incorrectly) characterize the harm as intended; and of course intended harm is almost always unjustifiable.¹⁶⁷ And, when harms are “on screen” and highly salient, psychological research teaches that people will feel that

¹⁶⁴ Some jurisdictions, indeed, define this exception so narrowly that only purpose, not individualized knowledge in any form, satisfies this exception. =cite from R.3rd. On the other hand, many jurisdictions treat deaths arising from Russian Roulette as “intentional” and not “accidental,” even if there is credible testimony that the insured decedent honestly believed that the live bullet was not in the firing chamber, and thus lacked (what I call) individualized knowledge of the risk. See, e.g., *Arnold v. Metropolitan Life Ins. Co.*, 970 F.2d 360, 362-363 (7th Cir. 1992).

¹⁶⁵ Useful discussions include: M. Adler & E. Posner, *New Foundations of Cost Benefit Analysis* (2006); *Cost-Benefit Analysis: Legal, Economic, and Philosophical Perspectives* (M. Adler & E. Posner 2002); *Risk: Philosophical Perspectives* (T. Lewens ed. 2007).

¹⁶⁶ If marketing a product in the face of a cost-benefit analysis predicting serious injuries or deaths from the use of the product counts as individualized knowledge, it is quite unlikely that existing legal doctrine would recognize a justification. As noted earlier, many jurisdictions decline to recognize a necessity defense for murder, or if they do recognize it, they limit the defense to “emergencies.” *Supra* note =. Securing benefits for consumers of a product (such as lower cost or greater ease of use) is not likely to count as a “necessity” justifying a knowing killing. Heinzerling, *Knowing Killing*, *supra* note 10= at 529.

¹⁶⁷ See TAN = *supra* (Knope on confusion of knowledge with intent).

actors have an especially strong duty to avoid those harms.¹⁶⁸ However, as we have seen, this psychological fact is too thin a reed to support a more stringent moral and legal duty.

Third, it is indeed plausible to critique some versions of cost-benefit analysis as morally and legally objectionable. But even when this is so, it is not at all plausible that the statistical knowledge of harmful effects that the actor acquires in the course of performing cost-benefit analysis is, by itself, what makes acting on the basis of that analysis objectionable. Rather, when such an analysis is objectionable, it is for different reasons—e.g., because distributive effects or deontological constraints are improperly ignored, or because the social welfare function gives inadequate weight to the value of human life or to soft variables that are difficult to measure.

Consider, for example, the special stringency principle and the examples of Alfa and Benna. Alfa knowingly creates close to a 100% risk of killing one pedestrian in order to save the lives of five passengers. Many deontologists would condemn her conduct, while many consequentialists would not. But the condemnation by deontologists does not flow simply from the fact that she can predict that her actions will cause harm. For they will not necessarily condemn all actors who possess only statistical knowledge that they will cause harm (recall Agatha and Larry). Indeed, they might not condemn the conduct of Benna, who knowingly creates a 20% risk of killing a pedestrian in order to save one passenger. Card-carrying deontologists can still endorse the need to balance consequences as part of the proper moral analysis, especially when the risk levels are smaller and are fairly distributed.¹⁶⁹

Fourth, it is easy to confuse the cold-blooded, deliberate state of mind that characterizes first-degree murder with the deliberate decision of an actor (such as a corporate manager) to move forward with a project or activity despite knowledge of its very likely harmful effects. By trading off health and safety against more diffuse social benefits, the actor who engages in cost-benefit analysis seems to be equally “cold-blooded.”¹⁷⁰ But this conflation of premeditated murder with “premeditated” use of cost-benefit analysis is unwarranted. The two categories are not directly comparable. To be sure, if D1 impulsively shoots another out of jealousy, while D2 deliberates over whether to shoot his victim out of jealousy and ultimately decides to do so, we might have reason to punish D2 more harshly (for first degree rather than second degree murder). Impulsiveness sometimes mitigates, and premeditation sometimes aggravates, in this context. But if corporate manager M1 impulsively decides to send a product onto the marketplace without conducting any advanced research into its benefits and risks, while M2 deliberately weighs those benefits and

¹⁶⁸ =cite. Recall the quotation from Sunstein, *supra* at = (comparing Company A with Company B). This is another instance of salience and the Availability heuristic. See *supra* at = (discussing these issues in the context of the identifiable victim effect).

¹⁶⁹ See Simons, *Tort Negligence*, *supra* note =, at =.

¹⁷⁰ See Heinzerling, *Knowing Killing*, *supra* note 10= at 523, 526. To be sure, if a product manufacturer can confidently predict that his product will contain some manufacturing defects that will cause harm, that prediction is sometimes relied upon in justifying strict liability for the resulting harms. See RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2 cmt. a (1998) (“Because manufacturers invest in quality control at consciously chosen levels, their knowledge that a predictable number of flawed products will enter the marketplace entails an element of deliberation about the amount of injury that will result from their activity.”). Unfortunately, this rationale is quite open-ended, justifying much broader strict liability than tort law actually imposes, or than tort law justifiably should impose. I do not have space to address the issue here.

risks, it hardly follows that M2 is more culpable. To the contrary, it is quite possible that M1 is negligent (or worse), while M2 is not culpable at all.

Fifth, even when an actor has conducted a cost-benefit analysis in an improper manner, the question remains: how culpable is this? The answer depends in part on what was defective about her analysis. A misunderstanding of the facts? A failure to conduct sufficient research? A failure to consider (or to give sufficient weight to) certain factors? But we should not assume that the decisive factor in determining her culpability for conducting an improper analysis is always the extent or severity of the harms of which she had statistical knowledge. Nor should we assume that a complete failure to conduct any cost-benefit analysis must be less culpable than undertaking such an analysis but performing it poorly. Sometimes, it is more culpable for a corporation (1) to fail, because of indifference, even to investigate the likely consequences of its decision to market a product that could produce serious harm if it is unsafe, than (2) to conduct a cost-benefit study of those consequences but to perform the analysis negligently.¹⁷¹ Or, in doctrinal terms, sometimes “recklessness” (in the sense of awareness of the predictable harms from the product) can be less culpable than “negligence” (in the sense of blameworthy inadvertence to whether the product will produce such harms). “In many respects the actor who does not care about the risks is worse than one that assesses the risks and makes the wrong judgment.”¹⁷²

The broader point, that cognitive recklessness is sometimes less culpable than negligence, raises some fundamental issues about how we should rank different mental states and different culpability categories.¹⁷³ But for our purposes, the lesson

¹⁷¹ Note that Viscusi’s study, mentioned in the introduction, *supra* at note 12=, provides empirical evidence that many people do not appreciate this point. In mock juror studies of punitive damages, participants tend to treat (2) as warranting much greater punitive damages than (1).

¹⁷² George Fletcher, *The Grammar of Criminal Law: American, Comparative, and International (Volume One: Foundations)* 315 (2007). Here is Fletcher’s argument in more detail:

Compare two possible versions of the prosecution against Ford Motor Company for locating the gas tank in a vulnerable position in their rear-engine model, the Pinto. In case A, the company never bothers to study whether it is safe or dangerous to build their car with gas tanks in front. In case B, the company undertakes a systematic study of the costs and benefits of expected deaths from collision and the costs of placing the tank in [a] safer spot. It concludes that in light of the difference in expected deaths from collision and the costs of redesigning the car, the saving of lives is not worth the costs. [Suppose in case B the costs of the design outweigh the benefits and the company is properly found to be reckless.]

Now suppose that the company was sloppy, it did not invest the money to test the design, and it was blissfully unaware of the risk. Do we seriously believe that a merely negligent Ford Motor Company, one that took exactly the same risks in marketing the Pinto, should receive more favorable legal treatment? The opposite conclusion seems plausible as well. In many respects the actor who does not care about the risks is worse than one that assesses the risks and makes the wrong judgment.

Id. at 314-315.

¹⁷³ In my view, acting with a cognitive mental state is sometimes more, and sometimes less, culpable than acting with a conative mental state (of intention or indifference) or than acting unjustifiably as measured by an objective standard of conduct. The conventional hierarchy of mental states, ranging from purpose, to knowledge, to recklessness, to negligence, is thus imperfect. See generally Kenneth W. Simons, *Rethinking Mental States*, *supra* note 15=.

is straightforward: a bare decision to conduct a cost-benefit analysis should not make the actor legally culpable, and even if such an analysis was conducted imperfectly or unreasonably, this should not inevitably make the actor legally culpable to the same degree as an actor who knowingly causes a harm.

Sixth, and finally, consider high-profile product liability cases in which the corporation's decision to trade off health and safety against the cost of a precaution apparently enraged the jury and prompted an award of punitive damages.¹⁷⁴ Such cases appear to be counterexamples to the arguments I have just offered. The best-known example here is the controversy over the alleged defect in the design of the gas tank of the Ford Pinto. Ford was found liable for an enormous punitive damage award, and was even criminally prosecuted. And a popular understanding of the reason for the award and for the prosecution is that Ford made a morally abhorrent, coldly calculated decision to weigh the cost of paying tort liability awards for predicted injuries and deaths against the lost profits it would suffer if it had to incur the cost (which was very small per automobile) of taking a precaution to prevent those harms.

This popular understanding rests on a contestable understanding of the facts. More importantly for our purposes, it is also unjustified as a matter of principle. With respect to the facts, it is not clear that the jury which imposed the punitive award, and the judges who upheld that award, decided as they did simply because Ford (a) performed a cost-benefit analysis of whether to make the gas tank safer and (b) thus improperly traded lives off against money,¹⁷⁵ though I concede that this simple abhorrence of cost-benefit balancing might indeed have been the central motivation.¹⁷⁶ However, if the popular understanding of the facts and of the

¹⁷⁴ For a recitation of some such cases, see W. Kip Viscusi, *Corporate Risk Analysis*, supra note 12, at =.

¹⁷⁵ See Schwartz, *id.* at 1039 (noting that the civil jury in *Grimshaw* was instructed under the consumer expectation, not a risk-benefit test); & 1059 n. 178 (noting that the enormous size of punitive damage award might have been based on Ford's profits over the relevant time period).

Many have also criticized Ford's decision, in its cost-benefit analysis, to value life at only \$200,000. However, that figure was supplied to Ford by a government agency. See *id.* at 1022-1026. Note also that the famous "smoking gun" cost-benefit memo in which Ford purportedly memorialized its cost-benefit analysis was not admitted into evidence. *Id.* at 1021.

¹⁷⁶ See *id.* at 1067. In *Grimshaw v. Ford Motor Co.*, 119 Cal. App. 3d 757 (Cal. App. 1981), the appellate court upheld the punitive damage award against Ford because a jury could find that it sold an automobile that it knew, based on crash tests, was highly dangerous:

Through the results of the crash tests Ford knew that the Pinto's fuel tank and rear structure would expose consumers to serious injury or death in a 20- to 30-mile-per-hour collision. There was evidence that Ford could have corrected the hazardous design defects at minimal cost but decided to defer correction of the shortcomings by engaging in a cost-benefit analysis balancing human lives and limbs against corporate profits. Ford's institutional mentality was shown to be one of callous indifference to public safety. There was substantial evidence that Ford's conduct constituted "conscious disregard" of the probability of injury to members of the consuming public.

...

While much of the evidence was necessarily circumstantial, there was substantial evidence from which the jury could reasonably find that Ford's management decided to proceed with the production of the Pinto with knowledge of test results revealing design defects which rendered the fuel tank extremely vulnerable on rear

motivation of the jury and appellate court were correct, and if Ford really was sanctioned with punitive damages simply for conducting a cost-benefit analysis, without regard to how sensitively and carefully they performed that analysis, then the punitive damage award was indeed unjustified.

Similarly, the 1980 criminal trial of Ford Motor Company for the same alleged design defect in the Pinto gas tank might appear to raise serious doubts about my argument; but it does not. Again, Ford was not prosecuted simply for proceeding with a product that Ford (statistically) knew would cause harm. Indeed, the State of Indiana prosecuted Ford for reckless homicide, not for murder; and yet murder (which under Indiana law encompasses knowingly causing a death) would have been an appropriate charge if indeed the law treats statistical knowledge the same as individualized knowledge.¹⁷⁷

C. Can retributivists tolerate the predicted mistaken punishment of the innocent?

Another highly instructive illustration of the confusions created by the statistical knowledge problem is an important debate about the justification of punishment. Consequentialist critics of retributivism charge that retributivists are inconsistent or hypocritical: according to the critics, retributivists cannot endorse as justifiable any humanly realizable punishment system, because the architects of any real-world system will know to a statistical certainty that the system will result in the mistaken punishment of some number of innocent defendants. This criticism fails, however, due to an improper conflation of statistical with individualized knowledge.

The basic debate has proceeded as follows. First, retributivists assert that consequentialist accounts of punishment (like consequentialist accounts generally) improperly permit social welfare to trump or override individual rights and principles of desert. A classic example of this problem, in the view of retributivists, is that consequentialists are committed to supporting the intentional punishment of identified innocents if this would serve a greater good—for example, they would support the lynching of an innocent if this were necessary in order to appease a mob and prevent much greater violence. Second, and in response, consequentialists suggest that the retributivists are not on such high moral ground, because every real-world and every practicable system of criminal justice will, unfortunately, make at least a few mistakes in attempting to identify the guilty. Retributivists are then faced with a dilemma: either they purport to justify such systems of punishment, even though we know for certain in advance that the result will be the mistaken punishment of the innocent; or they take the route of moral purity and irrelevance, and concede that no real-world system of justice is morally justifiable. If they choose the first horn, however, then they seem to be guilty of a fundamental

impact at low speeds and endangered the safety and lives of the occupants. Such conduct constitutes corporate malice.

Id. at 91-92 (underlying added). However, this excerpt is also susceptible to the more radical interpretation that any use of cost-benefit analysis in designing products is unjustifiable and warrants punitive damages. This view should be rejected, for reasons that I have explained.

¹⁷⁷ Indiana homicide law applicable to the 1978 crash defined murder as including knowingly causing death. Ind. Crim. Code 35-42-1-1. (The relevant provision was enacted in 1976).

According to Gary Schwartz, the jury's acquittal of Ford for the reckless homicide charge was due in part to the fact that the reckless homicide statute was only enacted in 1977, while Ford had manufactured the Pinto involved in the accident in 1973; thus, the theory of defect had to depend on Ford's failure to repair or warn of the defect, not on Ford's initial defective design. Schwartz, *supra* note 11=, 43 *Rutg. L. Rev.* at 1017.

inconsistency. On the one hand, they condemn consequentialists for sometimes permitting the greater good to justify the intentional punishment of the innocent. But on the other hand, in claiming to be able to legitimate a criminal justice system, they must rely on the very argument they condemn, for they rely on the overriding importance of a greater good (protecting the rights of victims or giving offenders their just deserts) in order to justify a lesser evil, the use of a punishment system that we know will inevitably punish a small number of innocents.¹⁷⁸

The analysis above helps to explain why this criticism of retributivism is too facile: it implicitly rests on an unfounded conflation of individual knowledge and statistical knowledge.

Suppose a judge decides, in an individual case, to punish a defendant he knows is factually innocent. This is both immoral and illegal, even if the judge reasonably believes that such punishment will provide substantial social benefits.¹⁷⁹ Such an action is indeed analogous to a judge deciding, in order to prevent mob violence, to free a known innocent so that private citizens can lynch him. But when a judge decides to participate in a system that he knows will result in some innocents being convicted (because he applies a reasonable doubt standard under which a very small number of innocent persons inevitably will still be convicted¹⁸⁰), his action is neither illegal nor immoral. And the difference is based on the distinction between individualized and statistical knowledge.¹⁸¹

¹⁷⁸ See Russell Christopher, *Deterring Retributivism: The Injustice of "Just Punishment,"* 96 *Nw. U. L. Rev.* =, 899-90= (2002).

Some try to rebut the criticism of retributivism by invoking the doctrine of double effect, distinguishing purpose (in the individual case) from knowledge (in a large number of cases). If an actor supports the lynching of an innocent in order to appease a mob, he intends to punish the innocent victim; but if a judge participates in a system that predictably results in mistaken convictions of the innocent, at worst he knowingly punishes the innocent but does not intend to do so. =; Moore, *Placing Blame* =, response to Dolinko?; Duff, recent draft =. But this move fails. Even if in an individual case the judge knowingly but not purposely punishes an innocent, his action is almost impossible to justify. And in this type of example, it is perfectly plausible to conclude that the judge acts with individualized knowledge, not with purpose. Thus, suppose that a judge finds a particular defendant guilty in order to protect the public, despite knowledge that that defendant is almost certainly innocent. If the defendant's innocence or guilt played no part in the judge's practical reasoning, it is more accurate to characterize the judge as knowingly but not purposely punishing an innocent.

¹⁷⁹ Suppose the judge had earlier presided over another trial involving this defendant in which he was acquitted despite overwhelming evidence of guilt; and suppose the judge reasonably believes that the defendant is a continuing serious threat to public safety.

¹⁸⁰ Many alternative systems are conceivable in which error might be almost entirely eliminated. Suppose conviction requires proof "beyond any doubt whatsoever," and requires two full trials before different juries.

¹⁸¹ The relevant distinction is not between "known" and "unknown" victims. See Dolinko, =, at 1633; Lempert =. Rather, it is between individualized knowledge in a single trial that the defendant is innocent, and statistical knowledge that across a large number of cases, some defendants are innocent. See, e.g., Larry Alexander, *Retributivism and the Inadvertent Punishment of the Innocent*, 2 *Law & Phil.* 233, 245 (1983) (critiquing the equation of "knowingly risking punishment of the innocent with knowingly punishing the innocent"). See the discussion of identifiability, *supra* at =.

In virtually¹⁸² no case is it morally justifiable to impose criminal punishment on an individual criminal defendant who the legal authorities know is innocent of the offense.¹⁸³ But, for the reasons we have canvassed above, it hardly follows that it is morally (or legally) impermissible for a judge to convict a defendant despite awareness of a genuine (but very small) statistical chance that he is innocent. Indeed, a judge who regularly presides over large numbers of criminal cases probably knows to a statistical certainty that he has imposed punishment on some innocent defendants. Rather, if it is morally and legally permissible to apply a reasonable doubt standard in a single case, given the relative error costs and the substantial social benefit of securing convictions of those who are guilty, then, under the principle of ICAA, it is also permissible to apply that standard in an indefinitely large number of cases. (And none of the exceptions to, or qualifications of, ICAA apply in this context.)

Now consider a scenario in which the actor with statistical knowledge is somewhat culpable. Suppose a judge participates in a system that, she realizes, creates an excessive, unjustifiable risk of convicting the innocent. (For example, she employs a jury instruction that is insufficiently stringent, one that she knows will often produce convictions even when the jury believes there is as much as a 10% risk that the defendant is innocent.) Although her participation in this system is wrongful (morally and legally), it is still less wrongful than upholding a conviction of an individual defendant who she knows is very likely innocent.

I conclude with two final observations about the problem of punishment of the innocent. First, even when legal actors possess neither purpose to punish an innocent person as such¹⁸⁴ nor individualized knowledge that the person being punished is innocent, their statistical knowledge that innocents are being punished is still morally and legally very significant, for it is highly relevant to their ongoing duty to minimize this unfortunate effect. A system that leads to this result demands a very compelling justification, and actors within the system have a stringent duty to minimize the risk of innocents being punished.

Indeed, a punishment system itself can be illegitimate when it permits too many innocent individuals to be convicted. Any punishment system in which 10% of those convicted were innocent simply creates too great a risk of error to be justifiable. And perhaps any capital punishment system that creates even a tiny risk

¹⁸² I include this qualification to recognize the possibility that where the national security is genuinely at stake, knowing or even purposeful punishment of an innocent might be warranted. But it is difficult to imagine an actual, credible example—even (or, given our experience, especially) after 9/11.

¹⁸³ I do not address whether legal innocence should be treated differently in this respect from factual innocence.

¹⁸⁴ The more accurate and honest way to conceptualize the “predictable punishment of the innocent” problem is as posing the question whether statistical knowledge should be treated in this context like individualized knowledge, and not as posing the question whether knowledge should be treated like purpose. For the special wrongness of punishing a known innocent individual remains even if the actor does not intend to be punishing an innocent person as such, i.e., even if the actor (a) intends to punish X, knowing that X is innocent, but (b) does not have as his purpose, or aim, or one of his reasons, that his acts punish an innocent person. Suppose the judge simply wants to ensure that X is punished, to avoid the adverse publicity from setting X free in light of the public’s strong but mistaken belief in X’s guilt; yet the judge hopes that X might actually be guilty, and certainly would prefer that X were not innocent and wishes he could acquit X and persuade the public of X’s innocence. The judge’s knowledge that he is punishing an innocent is still highly culpable.

of convicting the innocent is unjustifiable, given the irrevocability and seriousness of the death penalty.

In a recent, provocative article arguing that the death penalty is morally obligatory if it deters a significant number of private killings, Cass Sunstein and Adrian Vermeule explicitly claim that the risk of mistakes does not distinguish the death penalty from other social policies that predictably and inevitably cause or permit the deaths of innocents.¹⁸⁵ Many would reject their claim, believing that the statistically certain death of even one innocent due to the state's death penalty system is such a grievous wrong that no consequential benefits can justify it.¹⁸⁶ While I do not agree with Sunstein and Vermeule that legally authorized executions are on a moral par with construction projects and other human activities that predictably cause the death of innocent life, their insistence that human institutions inevitably produce some level of error is a valid point, worth remembering. And I am not convinced that the risk of even an infinitesimal risk of error should condemn a capital punishment system that is otherwise justifiable. On the other hand, there certainly is a great deal more we could do right now to reduce significantly the risk of erroneous executions in contemporary death penalty systems.¹⁸⁷

Second, the state should, of course, provide an ex post remedy if a case of mistaken punishment of the innocent is discovered. But it is a separate question whether the system that results in such mistakes is unjustifiable, and whether individual actors within the system bear moral or legal responsibility for such mistakes. The answer to these questions, I suggest, is “no”—or more precisely, “not necessarily.” We should not simply assume an affirmative answer based on an improper conflation of statistical with individualized knowledge.

D. Practical and doctrinal issues

I do not focus in this paper on practical and doctrinal issues. For example, significant difficulties of proof arise if legal sanctions depend on the specific mental state of an actor, including the precise probability of harm (or of the existence of a circumstance) that the actor believes to exist. Various legal strategies have been employed to circumvent these proof problems. For example, if individualized “knowledge” is the minimum mental state required, as it often is in criminal statutes, then judges and other legal actors may employ various devices to expand the meaning of “knowledge” in order to encompass lower-probability beliefs.¹⁸⁸ For purposes of this article, I simply assume that we have adequate proof of the cognitive mental state in question. But the proof difficulties can be serious, especially because we are considering whether the actor possesses a cognitive mental state. It is often easier to determine what an actor's purpose in acting is, than to determine what he knows about the circumstances of the act or its possible results, since his knowledge will often have no direct motivational significance.

¹⁸⁵ Cass R. Sunstein & Adrian Vermeule, *Is Capital Punishment Morally Required? Acts, Omissions, and Life-Life Tradeoffs*, 58 *Stan. L. Rev.* 703, (2006).

¹⁸⁶ Cf. Carol Steiker, *No, Capital Punishment is Not Morally Required: Deterrence, Deontology, and the Death Penalty*, 58 *Stan. L. Rev.* 751 (2005); Eric Blumenson, =.

¹⁸⁷ See Steiker, *id.*, mentioning inadequacy of trial counsel, =etc.

¹⁸⁸ Consider the doctrine of willful blindness, which essentially treats certain unusually culpable instances of reckless belief as satisfying a legal knowledge requirement. See Simons, *Retributivism Refined*, *supra* note 139=, at =.

Moreover, pragmatic concerns also shape the categories that the law employs: it must draw a somewhat arbitrary line between individualized knowledge and cognitive recklessness, identifying some particular probability formula (such as “practical certainty” or “highly probable” or “more likely than not”) for the former.¹⁸⁹

Although the analysis in the paper will certainly inform the answer to specific doctrinal questions, it will often not be dispositive. What counts as “individualized” rather than “statistical” knowledge might well depend on whether the issue arises in criminal law, tort law, or some other field of law, and indeed on whether it arises in one or another particular subcategory of criminal law, or of tort law. For example, whether repeated dangerous acts demonstrate a form of “defiance” that justifies punitive damages in tort is a different question than whether a repetition of similar acts justifies a recidivist premium in criminal law.

V. Conclusion

Acting with statistical knowledge is ordinarily less culpable than acting with individualized knowledge. Often, indeed, the first is not culpable at all. Expanding the spatial or temporal scope of an activity or a series of acts ordinarily does not increase the level of culpability properly attributable to the actor. The invariant culpability principles, ICAA and ICREA, formalize this idea.

We are now in a position to make sense of the various positions for and against distinguishing statistical from individualized knowledge that were set forth in the introduction, and also to show some of the implications of a more careful analysis of these concepts.

Consider first the arguments against distinguishing the two types of knowledge.

1. **“Statistical victims are still victims.”** True enough. The lives of all victims are indeed extraordinarily valuable, whether or not we can predict in advance the specific lives that will be lost. But the question before us is how best to measure the culpability of an actor who causes harm to victims; and in this context, we have reason to condemn more strongly (or at least to find some fault with) those actors who possess individualized, rather than statistical, knowledge that they will cause harm.¹⁹⁰ A primary explanation for this position is the special stringency principle (SSP), a deontological principle that treats an actor as highly culpable when he knowingly imposes a highly concentrated risk of serious harm on a victim.

2. **“Juries often reject the distinction between statistical and individualized knowledge.”** This behavior is a relevant datum. And indeed, rejecting the distinction is sometimes just what juries should do—specifically, when the risk of which the actor is statistically aware is extremely unjustifiable. However, often, juries simply make a mistake in rejecting the distinction, though their

¹⁸⁹ We would also need to draw a somewhat arbitrary line between cognitive recklessness and a lesser degree of cognitive culpability, if cognitive recklessness requires that the actor believe he is posing at least some specified threshold degree of risk, e.g. “substantial” or “significant” or “nontrivial.” See note = *supra* (noting that the Model Penal Code might impose such a threshold).

¹⁹⁰ By contrast, when we are considering *ex ante* what preventive measures government should take to prevent the death of innocents, the relative culpability of those who immediately cause such deaths might be a minor factor at best. It may be better to invest limited government funds to prevent 500 deaths caused by poor highway design than to prevent 20 murders. See TAN = *supra*.

behavior is understandable in light of cognitive biases, psychological heuristics, and a neglect of the invariant culpability principles.

3. **“Ordinary people often find cost-benefit analysis abhorrent.”** Yes, they do. Again, sometimes this is a justifiable reaction, as when the cost-benefit analysis is imperfectly done, or when the situation is one in which a deontological principle should trump or constrain the cost-benefit calculus (e.g. when SSP applies). But when people simply reject as morally objectionable all trade-offs between the advantages and disadvantages of taking precautions against risk, they express a moral position that is unrealistic in a world full of low-level risks and, more importantly, is untenable.

4. **“Acting with merely statistical knowledge is sometimes still highly culpable.”** This is correct, but only in limited circumstances—especially, when the risk is extremely unjustifiable (as in Time Bomb or Atmospheric Poison, but not in Atmospheric Cancer Cure).

Now consider the arguments in favor of distinguishing the two types of knowledge. With these assertions, I largely agree.

5. **“Scale should not matter.”** This is an important lesson of the invariant culpability principles, ICAA and ICREA. But we must also remember the *ceteris paribus* conditions that those principles contain. And we should recall that when a person’s individual act of risk-creation is faulty, repetition of that act might elevate the actor’s culpability to a higher level.

6. **“Careful risk/benefit analysis should not inculpate.”** Indeed it should not. This is another implication of the invariant culpability principles.

7. **“The modern hierarchy of mental states requires a distinction between statistical and individualized knowledge.”** True, though this argument only takes us so far, because in some respects the modern hierarchy is oversimplified. But in this context, it usually provides the right answer.

8. **“Knowingly harming another is distinctively wrongful only when the actor’s knowledge is individualized.”** This argument is indeed crucial, but that we must distinguish the following two versions of it.

(a) **“We owe a higher duty to identifiable victims.”** The term “identifiable” has many meanings, but none of them makes this claim persuasive, except insofar as the claim is a crude version of the following assertion.

(b) **“We owe a special duty not to impose highly concentrated risks of harm on a single person.”** Just so. And SSP is my attempt to formalize this idea. However, we should distinguish between a very high risk of serious physical harm to a person, and a very high risk that a legally relevant fact or circumstance exists—for example, the risk that an actor is in possession of stolen property or that the actor has sold alcohol to an underage person. It is doubtful that SSP applies in the latter context.

The analysis in the article has also suggested the following important conclusions.

- When a faulty actor repeats his unjustifiable acts or expands his activity, that repetition sometimes reveals the culpability of defiance of moral and legal norms. The invariant culpability principles, properly applied, recognize this special type of culpability. Accordingly, a retributivist can indeed support a punishment premium for recidivists.

- The legal criterion endorsed by the Restatement Third of Torts for distinguishing individualized from statistical knowledge is essentially sound, at least in cases where the knowledge concerns the causation of serious physical injury or death: individualized knowledge should normally embrace knowledge to a high likelihood that “the conduct will bring about harm to a particular victim, or to someone within a small class of potential victims within a localized area.”¹⁹¹ But the criterion must be qualified to account for unusual cases such as Programmed Rifle or Atmospheric Poison. In such cases, where the actor possesses merely statistical knowledge but his conduct is extremely unjustifiable, the actor’s culpability is comparable to that of an actor with individualized knowledge.
- A legal system can be legitimate even though legal actors within the system know that it will, as a statistical matter, punish the innocent. Retributivists are not, despite the claims of critics, vulnerable to the objection of hypocrisy or inconsistency for defending the legitimacy of an imperfect criminal justice system while condemning legal actors who punish an individual they know to be innocent. Individualized knowledge of the latter sort is much more culpable, and in need of a far more compelling justification, than statistical knowledge that the system will inevitably permit some innocent defendants to be mistakenly punished.

We have taken an extended journey through a varied “culpability” landscape populated with a motley cast of characters. The excursion did not fully resolve all the questions it explored. For example, I offered only a preliminary answer to the question why one might owe an especially stringent duty not to impose a concentrated risk of serious harm on an individual. Nevertheless, the article does provide a novel analytic structure, one that should assist future voyagers who wish to investigate this important, yet underexplored, piece of the terrain—the conundrum of statistical knowledge.

¹⁹¹ Note 139= *supra*.

VI. Appendix: formulas, principles, paradigms, variations

For the convenience of the reader, this appendix includes the formulas and principles discussed in the text, with some refinements. It also includes a series of paradigm individual knowledge scenarios, together with some variations.

Abbreviations:

IK = individualized knowledge

SK = statistical knowledge

A. Paradigm case of IK

Rolling Boulder 1:

Dee pushes a large boulder down the hill, for reason X, “knowing” (that is, believing to a probability exceeding T) that the boulder will strike and kill V.

In this paradigm case of IK, Dee’s conduct is presumptively highly culpable, warranting a severe legal sanction; and only an unusually compelling justification will render her conduct permissible. One explanation for these moral and legal implications is SSP (see below).

Implicit in the example, and in the variations that follow, is a definition of “knowledge.” The key part of the definition for our purposes is this: “knowing” that a result (such as V’s death) will occur requires that D must believe that the result will occur with a high probability, a probability exceeding some threshold T. What is the magnitude of T? In the law, T is usually understood as being greater than 50%, and often significantly greater. For example, in defining “knowledge,” the Model Penal Code requires “practical certainty” as to a result, or “high probability” as to a circumstance.¹⁹²

The example leaves open whether reason X helps justify the conduct or instead counts against justifying it. Thus, X could be a strong justification, e.g. Dee directs the boulder down the hill in order to plug a hole in a dike and thereby save 1000 people; or a strong anti-justification, e.g. in order to frighten Z at the bottom of the hill; or more neutral, e.g. for an idiosyncratic religious reason. The example also assumes that Dee does not push the boulder for the purpose of killing V.

B. Formulas and principles

1. ICAA = Invariant Culpability when Acts are Aggregated:

Suppose: (1) a person engages in a course of conduct containing multiple acts each of which, considered individually, is identical in all respects relevant to legal culpability, including:

the actor’s intentions, beliefs, and motives;

the socially recognized benefits and risks of harm that the act entails; and

the manner in which the benefits, and risks of harm, from the act are distributed among the actor, potential victims, and others;

¹⁹² =cites. See also Restatement Third of Torts, §1, comment c, requiring “substantial certainty” and indicating that this is not satisfied by “more probable than not.”

Then: (2) we should attribute to that actor the same degree of culpability for the whole course of conduct as the degree of culpability that we properly attribute to him for any individual act.

2. ICREA = Invariant Culpability when Risk-Exposures are Aggregated:¹⁹³

Suppose: (1) a person engages in *a single act that exposes multiple persons to a risk of harm, such that each risk-exposure is identical* in all respects relevant to legal culpability, including:

the actor's intentions, beliefs, and motives;

the socially recognized benefits and risks of harm that the act entails; and

the manner in which the benefits, and risks of harm, from the act are distributed among the actor, potential victims, and others;

Then: (2) we should attribute to that actor the same degree of culpability for the whole *set of risk-exposures* as the degree of culpability that we properly attribute to him for any individual *risk-exposure*.

3. SSP = Special stringency principle

If D knows that by his conduct he will impose a highly concentrated risk of serious harm on a single victim, then D has an especially stringent duty not to impose that risk, and is especially culpable if he does impose the risk. Consider this more precise, albeit still oversimplified,¹⁹⁴ formulation:

Action A creates an expected socially recognized benefit B and causes an expected socially recognized harm H.

$B = M_B \times P_B$ (where M_B is the magnitude of the social benefit and P_B is the probability that the act will secure that benefit)

$H = M_H \times P_H$ (where M_H is the magnitude of the social harm and P_H is the probability that the act will cause that harm)

SSP recognizes an "especially stringent" duty in this sense:

Even if a ratio n of B: H would (just barely) suffice to justify the conduct in the range of cases where P_H to an individual is relatively small, that ratio is insufficient to justify the conduct where P_H to an individual is highly concentrated, i.e., where $P_H > C$ and C is, say, 70%.

For an illustration, recall the Alfa/Benna example: even if Benna is justified in knowingly creating a 20% risk of killing one pedestrian in order to save the life of

¹⁹³ Differences from ICAA are *italicized*.

¹⁹⁴ A complete analysis should, among others things, consider all the socially relevant benefits and harms, not just one benefit and one harm.

one passenger, it does not follow that Alfa is justified in knowingly creating close to a 100% risk of killing one pedestrian in order to save the lives of five passengers.¹⁹⁵

What counts as a “highly concentrated” risk C triggering the rule? This could be pegged to T, the probability of risk perceived by the actor that counts as IK. But we might also understand SSP as a more continuous function, one not triggered by any particular threshold C or T.¹⁹⁶

C. Variations on paradigm case of IK

1. Paradigm case of IK

Rolling Boulder 1:

Dee pushes a large boulder down the hill, for reason X, knowing that the boulder will strike and kill V.

2. Variation: Consecutive acts endanger, and will likely kill, a single person

Rolling Boulder 2:

Dee pushes six smaller boulders down the hill, one after the other, for reason X,¹⁹⁷ knowing that each boulder endangers V. This time, she knows that each boulder poses a smaller risk (say, 20%) of killing V. However, the aggregate risk of death to V is still close to 100%.

Analysis: Dee knowingly imposes a highly concentrated risk on V, albeit through a series of acts. SSP applies. Her culpability is roughly comparable to that of Dee in Rolling Boulder 1.¹⁹⁸

3. Variation: Consecutive acts endanger six people, and will likely kill one of them

Rolling Boulder 3:

Dee pushes six smaller boulders down the hill, one after the other, for reason X, knowing that each boulder endangers a different person.¹⁹⁹ As

¹⁹⁵ Text at note = supra. Although the formulation in the text assumes a simple ratio function $n = (B/H)$, nothing turns on this. The function could be more complex, with n replaced by some new function $f(M_B, P_B, M_H, P_H)$. Whatever the function, the critical feature of SSP is that barely satisfying the new function in the range of cases where P_H is small, i.e. where P_H is much less than C, will be insufficient to justify the conduct in the range of cases where $P_H > C$.

¹⁹⁶ = [perhaps give further details about how this could work without a sudden threshold jump in special stringency].

¹⁹⁷ Suppose, for example, that successfully plugging the dike requires six boulders rather than one, or that frightening Z can be accomplished equally well either by rolling six smaller boulders or by rolling one large one.

¹⁹⁸ See “repeat-Benna,” a/k/a Benna*, supra note 83=.

These examples, as well as others in the appendix, assume for simplicity that the actor believes there is a 100% (or close to 100%) probability that the relevant harm would occur, either to an individual victim (by means of one act or many acts), or to one victim or another. But we can easily change the examples to ensure that they instantiate a requisite probability of 90%, or 75%, or whatever threshold T we want to employ for IK.

in Rolling Boulder 2, each act poses a small risk (say, 20%) of killing a victim. The aggregate risk that at least one of the six different people will die is still close to 100%.

Analysis: SSP would not normally apply where the risk to each victim is this small. And it should not apply just because the aggregate risk that at least one of a group of victims will die is close to 100%. After all, we could extend this variation to a case where D endangers 1,000 different victims, imposing a very small risk on each but an aggregate risk close to 100%; in such a case, clearly SSP should not apply. In the end, depending on X, Dee's conduct may or may not be justifiable; however, the features of the act identified in this variation are not sufficient by themselves to characterize Dee as acting with IK such that IK and SSP trigger the demand for an unusually strong justification.

4. Variation: One act, which will kill either A or B

Rolling Boulder 4:

Dee pushes a large boulder down the hill, for reason X. She also knows the following. Halfway down the hill, the path splits in two. The boulder will either travel down the left fork directly towards A or down the right fork directly towards B. It will almost certainly kill whoever is in its path. The probability of veering in either direction is 50%. Thus, A and B each has a 50% chance of being struck and killed.

5. Variation: One act, which will kill one of six persons

Rolling Boulder 5:

Dee pushes a large boulder down the hill, for reason X. She also knows the following. Halfway down the hill, the path splits into six forks. The boulder will either travel down the first fork directly towards A, or down the second fork directly towards B, or ... [and so on up to the sixth fork and sixth potential victim F.] The boulder will almost certainly kill whoever is in its path. The probability of its veering in each of the six directions is the same, i.e., 1/6 or 16.7%. Thus, A, B, C, D, E and F each has a 16.7% chance of being struck and killed.

Analysis of Rolling Boulder 4 and 5: These are difficult cases to classify.

From the victim's point of view, the risk is not highly concentrated, especially in Rolling Boulder 5, where the risk of death is only 16.7%. This would suggest that the SSP should not apply. But from the point of view of the actor, Dee, there is certainty that he will cause harm to a single victim, though he cannot tell in advance which person that will be. Is this enough to warrant the conclusion that he acted with IK, that he has knowingly killed whichever unlucky victim he kills?

I think not. These do not seem to be genuine examples of IK or of SSP applying. The impulse to aggregate these risks is understandable, especially if X is strongly unjustifiable. But how do we distinguish these cases from a more extreme case in which Dee's acts endanger 5,000 "downhill" possible victims, imposing a 1/5,000 risk of death on each? And how is this more extreme case distinguishable

¹⁹⁹ Assume in this and all the hypotheticals that the victim cannot escape the boulder. In this variation, perhaps different children will wander in the path, or perhaps Dee chooses six different boulder paths knowing that immobilized victims are at risk on each path.

from the situations of tunnel builders Agatha or Bertha? Once Agatha or Bertha “gets the ball (or boulder) rolling” and starts the huge tunnel project, she can confidently predict that physical harm to one or more workers will be a downhill side effect of her project.

D. IK with respect to a circumstance rather than a result

The Boulder example and variations all involved IK of a likely result of the actor’s conduct. We should also analyze cases where the actor has IK of the likely existence of a circumstance that renders his conduct illegal. Here, (1) SSP is much less likely to apply, but (2) other rationales for treating IK differently from SK do sometimes apply.

Consider three cases involving the crime of knowing receipt of stolen property, in which a pawn shop owner is found in possession of a stolen item.

Pawn Shop Owners P1, P2, and P3:

Owner P1 receives and offers for sale a particular piece of jewelry that he knows is stolen.

P2 adopts a policy for screening goods received by his pawn shop that he knows will result in 10% of the goods being stolen. (He requires any form of identification from persons selling goods to him.) P2 receives and offers for sale a particular piece of jewelry, not realizing it is stolen.

P3 adopts a more effective screening policy that he knows will result in only 1% of goods being stolen. (He requires a photo identification from persons selling goods to him and also sends a list of all newly received property to local police so that they can check it against their list of recently stolen property.) P3 receives and offers for sale a particular piece of jewelry, not realizing it is stolen.

Suppose we conclude that P2’s policy is negligent, while P3’s is reasonable.

Analysis: P1 clearly acts with IK. But P2 and P3 arguably do not. Rather, these are arguably SK cases.

In all three cases, the special concerns about not imposing highly concentrated risks on individual victims (SSP) do not seem to apply. To be sure, in all three, the actual jewelry owner is directly harmed when a pawn shop improperly takes possession of his property. But it is doubtful that an especially stringent deontological constraint applies when one creates, say, a 70% risk of improperly depriving a victim of her property but not when one creates a 30% or 10% risk.

However, the conclusion that SSP does not apply does not answer the further question: should we treat P2 (or even P3) as IK for other reasons? With both results and circumstances, there are reasons other than SSP for giving special legal significance to IK cases. Recall the aggravated culpability and “greater need for justification” rationales for the IK category. Do these apply to P2? The answer is unclear.²⁰⁰

²⁰⁰ Here are some considerations.

1. It certainly is the case that as the probability of harm (either of a harmful result or of a harmful circumstance) that the actor subjectively perceives increases, we have more reason to treat the actor as displaying aggravated culpability and to demand a more compelling justification before his action will be deemed lawful. Thus, interpreting a statutory “knowledge” requirement as punishing only P1 in the above example is defensible.

Now consider another set of examples that test my argument that SSP does not apply when the concentration of risk pertains to a circumstance rather than a result. The following is an especially sympathetic “circumstance” case for applying SSP. The example is highly stylized in order to present a helpful contrast with the previous example.

Actors S1, S2, and S3: sexual intercourse with a nonconsenting victim:

Actor S1 knows that his sexual partner V1 is not consenting,²⁰¹ but he nevertheless proceeds to have sexual intercourse with her.

S2 adopts a policy for determining whether his sexual partner consents that he knows will result in a 10% chance that his partner does not consent. (The “policy” is: he simply initiates and proceeds with the act of intercourse, making no effort to determine whether his partner consents; if she is passive or silent, he continues.) S2 has intercourse with V2, not realizing that she does not consent.

S3 adopts a more effective screening policy that he knows will result in a 1% chance that his partner does not consent. (The “policy” is: he carefully considers his partner’s facial expression and body language and if satisfied that she is genuinely willing, he proceeds with the act of intercourse; he does not specifically ask whether she consents.) S3 has intercourse with V3, not realizing that she does not consent.

Suppose we conclude that S2’s policy is negligent, while S3’s is reasonable.²⁰²

Analysis: Clearly S1 acts with IK. But again, S2 and S3 probably do not. Rather, they are SK cases. Of course, S2 can be punished if we expand the required culpability for sexual assault from IK to recklessness or to negligence.

In this example, however, a stronger argument exists that the special concerns about not concentrating risks on individual victims (SSP) does apply. For S1 has decided to impose on V1 a very high risk of infringing sexual autonomy, and this is arguably not just proportionately worse, but disproportionately worse, than S2 imposing a smaller risk of infringing the autonomy of V2, or S3 imposing an even smaller risk of that sort on V3. The culpability of S1 is arguably disproportionately greater than the culpability of S2, in much the same way that the culpability of Alfa is disproportionately greater than the culpability of Benna in the earlier example.

2. On the other hand, some context-specific features of a particular legal norm might suggest a broader interpretation of IK. In the receiving stolen property context, for example, it is increasingly more burdensome for a business owner to avoid causing or creating a risk of the criminal harm, as we move from P1 (the burden of not acquiring a particular item known to be stolen), to P2 (the burden of not creating a small risk that she is in possession of stolen property), to P3 (the burden of avoiding a very small risk). And perhaps the difference in burden is quite dramatic when we move from avoiding a small to avoiding a very small risk; so it makes some sense to impose legal liability under the “IK” rubric for P2 but not for P3.

3. And yet, on the third hand, it is much simpler for legal actors to employ a relatively consistent conception of IK across contexts; this factor suggests that we should restrict “IK” to P1 and, if we want to sanction P2, characterize P2 as a case of recklessness, not IK.

²⁰¹ Assume nonconsent means subjective unwillingness, whether or not expressed to the other.

²⁰² If this conclusion is doubtful with respect to S3 on these facts, we can sufficiently lower the “error rate” for S3 and alter his “strategy” in order to make the conclusion true.

E. Russian Roulette cases

These cases are especially vivid. They involve clearly unjustifiable conduct, so they lack the useful generality of the Rolling Boulder paradigm case and variations. But in many respects they are parallel to those cases. Assume that the actor in these examples is aware of the risks (a more plausible assumption here than in many other scenarios), and intends to expose the victim to risk, but does not intend to kill the victim.

Russian Roulette 1 (Malone) (multiple shots, single potential victim):

Malone places one bullet in a five-chamber revolver, spins the cylinder, and fires at the victim three times, causing death with the third shot.²⁰³ He knows that the cumulative risk of death is 60%, while the risk created by each separate shot is less.

Analysis: Malone's conduct is akin to Dee rolling six boulders towards V in Rolling Boulder 2. He imposed a 20% chance of death with the first shot, a 25% chance with the second, and a 33.3% chance with the third, and a cumulative risk of 60%. If IK requires a 60% probability of death, then Malone has IK; if it requires more, then he does not.²⁰⁴

Russian Roulette 2 (Stallone) (multiple shots, multiple potential victims):

Stallone places one bullet in a five-chamber revolver, spins the cylinder, and fires the gun once at A, then once at B, then once at C, causing C's death with the third shot. He knows that the cumulative risk that he will cause death either to A, to B, or to C is 60%, while the risk of death he has created to each distinct victim by each separate shot is less.

Analysis: We cannot fairly characterize Stallone as having killed C with IK, just as we cannot characterize the actor in Rolling Boulder 3 as having killed the victim with IK.

Once again, Stallone can still be convicted of murder if the requisite culpability for murder is not restricted to IK and purpose to kill, but also embraces depraved heart. If the culpability were restricted, it would have the unfortunate effect that an actor could deliberately avoid murder liability by restructuring his method of killing: instead of shooting a random person in the head, terrorist Stallone* could simply spin the cylinder containing one bullet and fire five times to the head of five different people; instead of exploding a bomb in a workplace where he knows who will die, he could explode it in a busy public square where he does not know; and so forth.²⁰⁵

²⁰³ Adapted from *Commonwealth v. Malone*, 47 A.2d 445 (Pa. 1946).

²⁰⁴ Of course, an alternative ground for murder liability is the ground that the court in Malone endorsed: depraved heart murder, which does not require IK.

²⁰⁵ Leo Katz has carefully explored the issue of the amenability of deontological constraints to certain forms of manipulation. See Leo Katz, *Ill-Gotten Gains: Evasion, Blackmail, Fraud and Kindred Puzzles of the Law* (1996); Katz, *Before and After: Temporal Anomalies in Legal Doctrine*, 151 U. Pa. L. Rev. 863, 879-880 (2003) (comparing two characters analogous to my Maxwell and Ned).