ARTICLES

STATISTICAL KNOWLEDGE DECONSTRUCTED

KENNETH W. SIMONS[*]

The law frequently distinguishes between individualized knowledge, or awareness that one's act will harm a particular victim, and statistical knowledge, or awareness that one's activity or multiple acts will, to a high statistical likelihood, harm one or more persons from a large class of potential victims. (Compare driving through an intersection while aware that one's automobile is likely to injure a pedestrian in the crosswalk with managing 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 remarkably difficult to explain and defend.

This Article – the first systematic analysis of this pervasive but underappreciated problem – offers a qualified defense of the distinction between individualized and statistical knowledge. Acting with statistical knowledge is ordinarily less culpable than acting with individualized knowledge, and it is often 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 expansion 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, it is much harder to justify (a) speeding to the hospital to save five passengers, with knowledge that this will very likely require killing a pedestrian in one's path, than (b) speeding to the hospital to save one passenger, with knowledge 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.*

- *If 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 in the face of such knowledge does not, by itself, demonstrate that the activity is unjustifiable or that he is culpable for so proceeding.*

- *A legal system can be legitimate even though legal actors within the system know that it will, as a statistical matter, punish the innocent.*

- *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.*

INTRODUCTION

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

Clara (impatient driver acting with individualized knowledge)

Clara is running late for a flight to a critically important business meeting. Driving along a one-lane deserted mountain road, she suddenly sees a man lying on the road in the path of the car. She honks. He stirs but does not move out of the way. Unable to drive around him and unwilling to stop the car, she slows down and drives over his body. She realizes that she is very likely to injure him severely, which she does.[1]

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

But the Anglo-American legal system does not treat actors so severely in cases where the knowledge is "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.

---

[1] The illustration is loosely based on a famous hypothetical from Sandy Kadish's criminal law casebook. *See* SANFORD H. KADISH, STEPHEN J. SCHULHOFER & CAROL S. STEIKER, CRIMINAL LAW AND ITS PROCESSES 842 (8th ed. 2007).

[2] Other discussions of the problem also refer to "statistical" knowledge. *See* Jody Armour, *Interpretive Construction, Systemic Consistency, and Criterial Norms in Tort Law*,

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

Agatha (careful builder acting 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 during the ten-year period of the project will 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, and 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? Is she liable for punitive damages? Knowingly causing a death usually counts as murder. Is she guilty of the murder of the worker who died? Under current law, the answer to all of these questions is no.

Now consider this variation:

Bertha (less careful builder acting with statistical knowledge)

Bertha is building an enormous tunnel, of the same size and scope as Agatha's. Bertha takes less extensive precautions than are customary, based on advice that these lesser precautions are almost as effective but are less costly and less 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.

---

54 VAND. L. REV. 1157, 1158-59 (2001); James A. Henderson & Aaron D. Twerski, *Intent and Recklessness in Tort: The Practical Craft of Restating Law*, 54 VAND. L. REV. 1333, 1141-42 (2001); Sanford H. Kadish, *Respect for Life and Regard for Rights in the Criminal Law*, 64 CALIF. L. REV. 871, 891-94 (1976). Thomas Schelling appears to have originated the terms "statistical life" and "statistical death," to be contrasted with losses of identified, individual lives. *See* Gregory C. Keating, *Pricelessness and Life: An Essay for Guido Calabresi*, 64 MD. L. REV. 159, 160 n.4 (2006) (citing THOMAS SCHELLING, *The Life You Save May Be Your Own*, *in* CHOICE AND CONSEQUENCE 115 (1984)).

[3] For simplicity, assume that she is the owner of the construction company and makes all relevant decisions about taking precautions. This obviates the difficulty 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 such laws do not bar this particular category of workers from tort recovery.

Although Bertha might be found negligent, or perhaps reckless, in deciding to use lesser precautions, it is quite unlikely that a court will find her 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? Conversely, if in the second scenario we treat Bertha as a "knowing" injurer or killer, why should we not so treat Agatha in the first scenario, even though, by hypothesis, Agatha used all reasonably feasible care? Moreover, recall Clara, the impatient 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 Kadish[4] that vividly highlights the contrast between actors with individualized knowledge and actors with statistical knowledge:

Eleanor (builder acting 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; 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

---

[4] *See* 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. *Id.*

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

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 reinforces 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.

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

- ***Scale alone should not affect culpability.*** If you drive carefully on a single occasion, imposing low risks on others, you do not deserve legal sanction if you happen to cause harm. Why, then, should you suddenly become liable merely because you choose to drive in the same manner every day? If you hire one worker after carefully checking his documentation, you do not deserve legal sanction if it turns out that he is an undocumented immigrant. Why, then, should you suddenly become liable for knowingly hiring an undocumented immigrant[6] if you employ 100,000 workers, and know that your careful screening will nevertheless result in your hiring at least one undocumented worker?

  Below, I formalize the argument that multiplying similar acts or similar occurrences of risk-exposure, either over time or space, does not by itself change the culpability of the underlying act. The "Invariant Culpability when Acts are Aggregated" principle (ICAA) applies whenever similar acts are multiplied, and the "Invariant Culpability when Risk-Exposures are Aggregated" principle (ICREA) applies when similar occurrences of risk-exposure are multiplied.[7]

- ***Careful risk/benefit analysis should not inculpate.*** Many believe that sound and justifiable regulatory and tort law policy requires companies to conduct careful risk-benefit analyses of whether to adopt precautions that would reduce the health and safety risks of their products or activities. Imposing sanctions on companies merely because they conducted such analyses and carefully examined the probable consequences of their actions would be inconsistent with this

---

[5] Kadish, *supra* note 2, 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. § 1324(a)(1)(A) (2006).

[7] *See infra* Part II.C, II.E.1.

desideratum. If statistical knowledge should be treated as harshly as individualized knowledge, however, it follows that if a company chooses to engage in far-flung activities or to manufacture or distribute a large number of products, it should be treated as highly culpable simply because it conducted a careful cost-benefit analysis and the scale of its activities is vast. That is a most unwelcome conclusion – and a conclusion explicitly rejected by language in the recently adopted *Restatement (Third) of Torts*.[8]

- ***The modern hierarchy of mental states requires a distinction between statistical and individualized knowledge.*** Under 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.[9] But if statistical knowledge is not distinguished from individual knowledge, this framework crumbles, because many cases of statistical knowledge seem indistinguishable from cases of recklessness, negligence, or even faultless conduct.

- ***Knowingly harming another is especially 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.* If the actor has statistical knowledge, the victim is typically unidentified, while if he has individualized knowledge, the victim is identifiable. Arguably, we have a more stringent moral and legal obligation to save an identifiable victim than to take advanced precautions against the risk of harm to unidentified possible future victims. Governments and coal

---

[8] RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL OR EMOTIONAL HARM § 1 cmt. e (2010).

> [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.

Id. For further discussion of the *Restatement (Third)*'s position, see *infra* Part II.E.1.

[9] *See* MODEL PENAL CODE § 2.02 (1985).

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.

(2) *We owe a special duty not to impose highly concentrated risks of harm on a single person.* On this view, concentrating the costs of an activity on one person merely to provide small benefits to many other persons is wrongful. 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 Gregory 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 show . . . .[10]

The objection is powerful and persuasive. Clearly, it would be wrong to knowingly permit the technician's death simply to provide a small benefit to each viewer, 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.

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 a mere statistic, not identifiable in

---

[10] Gregory Keating, *Pressing Precaution Beyond the Point of Cost-Justification*, 56 VAND. L. REV. 653, 666 (2003). I have changed Professor Keating's original example from a broadcast of a popular television show to a broadcast of the World Cup.

advance, then the victim's death somehow counts for less. But as Professor Lisa Heinzerling has argued, "Simply 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."[11] And arguably it is immoral, 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."[12]

- *Juries often reject the distinction.* Juries often look harshly upon a company's deliberate decision not to adopt a safety precaution when the 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 challenging Ford's decision not to strengthen the Pinto's fuel tank and General Motors's decision 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 the 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.[13]

---

[11] Lisa Heinzerling, *Knowing Killing and Environmental Law*, 14 N.Y.U. ENVTL. L.J. 521, 531 (2006) [hereinafter Heinzerling, *Knowing Killing*]; *see also* GUIDO CALABRESI, IDEALS, BELIEFS, ATTITUDES, AND THE LAW 6 (1985) ("I have never run into a 'statistical' life anywhere, in my car or anyplace else."); Lisa Heinzerling, *The Rights of Statistical People*, 24 HARV. ENVTL. L. REV. 189, 190 (2000) [hereinafter Heinzerling, *Statistical People*].

[12] *See* James K. Hammitt & Nicolas Treich, *Statistical vs. Identified Lives in Benefit-Cost Analysis*, 35 J. RISK & UNCERTAINTY 45, 45 (2007) (defining statistical murder as the preference of modern societies to prevent death to identified victims rather than anonymous victims (citing Joanne Linnerooth, *Murdering Statistical Lives . . . ?*, in THE VALUE OF LIFE AND SAFETY 229, 230-31 (M.W. Jones-Lee ed., 1982))); *see also* Allan Gibbard, *Risk and Value*, in VALUES AT RISK 94, 97 (Douglas MacLean ed., 1986) (wondering why building and operating a construction plan that is statistically certain to kill people is not "murder for hire").

[13] In the 1978 Ford Pinto case, the initial punitive damage award was $125 million (an extraordinarily large sum at that time), which was later reduced to $3.5 million. W. Kip Viscusi, *Corporate Risk Analysis: A Reckless Act?*, 52 STAN. L. REV. 547, 569 (2000). 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 Y. Hong, *Judge Cuts Award Against GM to $1.2 Billion*, L.A. TIMES, Aug. 27, 1999, at B1; *GM to Settle Case Over Gas Tank Explosion*, L.A. TIMES, July 25, 2003, at 4. *See generally* Gary Schwartz, *The Myth of the Ford Pinto Case*, 43 RUTGERS L. REV. 1013 (1991).

- ***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.[14]

  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.[15] These studies
  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 – 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 nearly 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 the following example:

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.

---

[14] *See* Viscusi, *supra* note 13, at 552-58.

[15] *See id.* at 588. Although Viscusi's study is illuminating, one 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 J. MacCoun, *The Costs and Benefits of Letting Juries Punish Corporations:
Comment on Viscusi*, 52 STAN. L. REV. 1821, 1823-24 (2000).

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? Yet it is indeterminate in advance whom the bomb will kill and when they will die.

This last example seems to suggest that other actors who possess statistical knowledge and 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"? Similarly for the tunnel builders Agatha and Bertha, isn't the decision to embark on a huge construction project akin to triggering a bomb or distributing a poison that will have destructive effects on someone at some point?

***

The conundrum of statistical knowledge 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 doctrinal and theoretical questions:

(a) What is the basic rationale for treating individualized knowledge as an especially culpable state of mind?

(b) Do recidivists deserve a punishment premium?

(c) Why do many people intuitively, but often wrongly, condemn corporate decision-makers for proceeding in the face of a cost-benefit analysis that predicts that their decisions not to take a safety precaution will cause death or serious harm?

(d) How can a legal system be legitimate, even though it predictably punishes the innocent?

(e) Do actors owe a more stringent duty not to harm "identifiable" victims?

The Article is organized as follows. Part I addresses the legal relevance of mental states in general, and cognitive mental states in particular. Part II.A explains the scope of the problem. Part II.B introduces the ICAA principle, "invariant culpability when acts are aggregated," as a novel analytical device, while Part II.C examines whether the principle applies to faulty as well as faultless actors. Part II.D elucidates SSP, the special stringency principle. Many nonconsequentialists 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. Part II.E considers some possible criteria for distinguishing statistical from individualized knowledge; delineates the principle of ICREA, invariant culpability when risk-

exposures, rather than acts, are aggregated; and explains that in exceptional circumstances, particularly when the risk is *extremely* unjustifiable, statistical and individual knowledge are indeed normatively indistinguishable. But these exceptions actually "prove" the rule, i.e., they actually underscore the legitimate rationales behind ICAA and ICREA.

Part III applies the analysis to several issues, including contexts where knowledge pertains to a legally relevant circumstance rather than a causal result, hostility to cost-benefit analysis, and a serious objection to retributivism. The objection is that retributivists cannot justify any real-world system of punishment, because they 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. The Conclusion responds to the questions I have raised in this Introduction.[16] An Appendix explores the relevance, if any, of various heuristic errors and

---

[16] I should briefly note some limitations and assumptions of this Article's analysis. First, in speaking of an actor's belief 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 D. Adler, *Against "Individual Risk": A Sympathetic Critique of Risk Assessment*, 153 U. PA. L. REV. 1121, 1226-27 (2005) [hereinafter Adler, *Against "Individual Risk"*]. A frequentist analysis of risk, however, might be relevant to certain issues, such as loss-of-a-chance damages in tort law. *See* Matthew D. Adler, *Risk, Death, and Harm*, 87 MINN. L. REV. 1293, 1436-44 (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 this 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 LOY. L.A. L. REV. 1171, 1218-21 (2008). 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 ten 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 nine deaths. D1 alone, however, creates an extremely high individual risk to the group he endangers, triggering what I call the Special Stringency Principle. *See infra* Part II.D. 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 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 to resolve the question (e.g., "a pedestrian was injured by an unidentifiable bus and 60% of the buses in this town are owned by Blue Bus Co."). *See* Charles Nesson, *The Evidence or the Event? On Judicial Proof and the Acceptability of Verdicts*, 98 HARV. L. REV. 1357, 1378-79 (1985). A valuable recent discussion of these evidentiary issues is presented in Mike Redmayne, *Exploring the Proof Paradoxes*, 14 LEGAL THEORY 281 (2008).

cognitive biases to the previous analysis, and it explains why the underlying issue in distinguishing individualized from statistical knowledge is not whether the victim is identifiable.

I.   THE LEGAL RELEVANCE OF COGNITIVE MENTAL STATES

As background to the analysis to follow, it is important to outline briefly some of the principal reasons why mental states in general, and cognitive mental states in particular, matter to legal liability. An actor's state of mind is frequently a crucial determinant of his legal responsibility. State of mind 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 and innocence, or between mandatory life imprisonment and a much shorter prison term. A tort defendant's mental state determines whether the defendant has committed an intentional tort or a tort of negligence. That distinction entails further consequences, such as 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," not if she is merely negligent.

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.[17] In this

---

[17] 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, belief, desire, justifiability, and deviation from a normative standard. Purpose or intention requires a "conscious object" to bring about a result (for example, murder requires a conscious object to cause the death of a victim) and thus typically requires both a belief that the result is possible and a desire that the result occur. MODEL PENAL CODE § 2.02(2)(a)(i). Knowledge, insofar as it is simply a mental state, requires a belief that a result is quite likely to occur. *Id.* §§ 2.02(2)(b), 2.02(7); RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL OR EMOTIONAL HARM § 1(a)-(b) (2010) (defining "intent" as acting with the purpose of causing a particular result or acting with the knowledge that the result is "substantially certain" to occur); RESTATEMENT (SECOND) OF TORTS § 8A (1965). Recklessness, in a common modern formulation, requires both a belief that there is a significant or substantial risk that the harmful result will occur, and that taking the risk is unjustifiable. MODEL PENAL CODE § 2.02(2)(c). And negligence is often defined as culpable inadvertence to risk; it requires a deviation, or gross deviation, from the standard of care that a reasonable person would observe in appreciating a risk or in acting to avoid or minimize the risk. *Id.* § 2.02(d).

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 fashion. 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.

Article, I focus on 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 colloquially speak of "knowledge" that one's conduct creates a *small risk* of harm,[18] or of "knowledge" that there is a *small risk* that a fact exists.[19] These examples, in which the actor is aware of a significant risk of harm rather than a high likelihood of harm, are instances of cognitive "recklessness," not knowledge, under modern doctrine.[20]

In principle, in judging the degree of an actor's responsibility, the law could employ numerous belief requirements spanning the entire range of probabilities, from a belief that the probability is 100% to a belief that it is just above zero. As a practical matter, 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 between awareness of a substantial risk, awareness of a high probability, and simple belief that the circumstance exists.[21] Some other criminal codes identify only one point, and tort law standards often identify only one.[22]

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. If Don is engrossed in reading Professor Simons's 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.

[18] For example, she "knew" that her dangerous driving might kill someone.

[19] For example, where the fact is the victim's lack of consent, the actor "knew" that the victim might not have consented to their sexual relations.

[20] Sometimes, however, recklessness criteria do not require subjective awareness of risk; instead, they require either gross negligence or some form of culpable indifference. *See* RESTATEMENT (THIRD) OF TORTS: LIAB. FOR EMOTIONAL OR PHYSICAL HARM § 2 cmt. c. 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. *See* MODEL PENAL CODE § 2.02(2)(d). 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 merely because it is based on negligent beliefs. *See* RESTATEMENT (THIRD) OF TORTS: LIAB. FOR EMOTIONAL OR PHYSICAL HARM § 3.

[21] *See* MODEL PENAL CODE § 2.02(2)(a)(ii), 2.02(2)(c), 2.02(7). The Model Penal Code

We should also remember that "recklessness," as defined in the influential Model Penal Code and in many state criminal codes, is a complex concept that combines 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.[23] In most of this Article, I will be using the term in its simpler, purely cognitive sense, i.e., as encompassing requirement (a) but not requirement (b).

What is the legal 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 (as 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 more culpable than acting with the "lower"

---

treats the awareness of a high probability that a material circumstance exists and the belief that the material circumstance exists 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." *Id.* § 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 ST. J. CRIM. L. 179, 182-83 (2003). With respect to results, the MPC defines knowledge as belief that the result is "practically certain." MODEL PENAL CODE § 2.02(2)(b)(ii).

[22] Under the *Restatement (Second) of Torts*, intent is defined to include both purpose and knowledge. *See* RESTATEMENT (SECOND) TORTS § 8A. Negligence does not require actual awareness of risk. *See id.* § 282 cmt. e. Recklessness is an infrequently used mental state in tort law, and it is not clear whether it requires such awareness. *See* Geoffrey Christopher Rapp, *The Wreckage of Recklessness*, 86 WASH. U. L. REV. 111, 147-48 (2008); Kenneth W. Simons, *Rethinking Mental States*, 72 B.U. L. REV. 463, 482-83 (1992). It is similarly unclear whether the reckless disregard often required for punitive damages, although sometimes defined differently from the recklessness occasionally used as a basic mental state for liability, requires awareness of risk. *See id.* at 484.

[23] MODEL PENAL CODE § 2.02(c) (defining as reckless a person who "consciously disregards" a risk that is "substantial and unjustifiable").

belief of cognitive recklessness. Not surprisingly, knowingly causing death is usually murder, while recklessly causing death is usually manslaughter.[24] 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 intentional torts, negligence, or 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*, more responsible or culpable than a person who creates a risk believing only that a harmful consequence is possible; the first person's action reveals willingness to tolerate a greater harm or evil in pursuit of his ends. When we condemn an actor for knowingly or recklessly causing harm, we are not condemning her simply for *possessing* a "guilty" or culpable state of mind as to the possible result. More precisely, we hold her responsible for her willingness to act notwithstanding her belief that her actions will or might cause harm.[25] What is morally and legally significant is her decision to act *in light of* her beliefs.

Second, consider the threshold function of a mental state. Knowledge requirements often serve this function. This is true not only of many crimes[26] 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.[27]

---

[24] 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. *See, e.g.*, State v. Aiwohi, 123 P.3d 1210, 1214 (Haw. 2005). 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.

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

[26] For example, knowledge is often required for drug and other possession offenses. *See, e.g.*, MODEL PENAL CODE § 2.01(4).

[27] Negligent invasions of those interests are not always actionable under the tort of negligence. The latter tort requires that the victim suffer physical harm; the torts of trespass

Third, consider the justificatory function of mental states. Actions accompanied by more "serious" or culpable mental states require more compelling justifications and often require narrower categories of justification, 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.[28] But an actor 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., taking a drive to feel the wind rushing through his hair), so long as the risk of death he poses is relatively small.

II.   STATISTICAL KNOWLEDGE DECONSTRUCTED

A.   *Specifying the Problem*

Recall tunnel builders Agatha and Bertha. Careful Agatha can 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 can predict that, with a lower, probably negligent, level of precaution, injuries and death are even more likely. And recall Clara, the impatient driver who knowingly injured someone lying in the road in order to catch her plane. 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.[29] Yet, in one sense,

---

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, 1079-80 (2006).

Similarly, as a constitutional matter, tort liability for defamation of a public figure or public official requires proof that the defendant at least "recklessly disregarded" the truth. *See* N.Y. Times Co. v. Sullivan, 376 U.S. 254, 279-80 (1964). That phrase has been interpreted as requiring a state of mind close to knowledge. *See* Saint Amant v. Thompson, 390 U.S. 727, 731 (1968).

[28] Indeed, many states do not permit a necessity defense to murder. *See* JOSHUA DRESSLER, UNDERSTANDING CRIMINAL LAW § 22.04[B], at 298 (5th ed. 2009). When they do recognize necessity, they often limit the necessity to "emergencies." *Cf. id.* § 22.05, at 300.

[29] There is little explicit case law on the topic. Some tort cases, however, have held that 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

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 not often noted, yet it poses analogous questions. Recall the earlier example of a large corporation charged with the crime of "knowingly" hiring undocumented immigrants as employees. 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.[30] But why, exactly, would that broad interpretation be improper?

---

reach some non-smokers, the Court finds that such generalized knowledge is insufficient to satisfy the intent requirement for battery."); Madden v. D.C. Transit Sys., Inc., 307 A.2d 756, 756-57 (D.C. 1973) (rejecting a battery claim for lack of requisite intent where the plaintiff claimed that the defendant knew that its buses were regularly spewing fumes and oily substances).

Some academics have supported tort liability for statistically known harms. *See, e.g.*, Paul A. LeBel, *Intent and Recklessness as Bases of Products Liability: One Step Back, Two Steps Forward*, 32 ALA. L. REV. 31, 67 (1980) (arguing that manufacturers' risk-utility balances are faulty because they do not consider the cost of public outrage over the use of such a balance and that imposing liability on an intentional tort theory would force manufacturers to internalize this cost); Darren S. Rimer, *Secondhand Smoke Damages: Extending a Cause of Action for Battery Against a Tobacco Manufacturer*, 2 SW. U. L. REV. 1237, 1240-41 (1995) (arguing that statistical knowledge should suffice for intentional tort liability for product defects).

[30] My research has not uncovered cases containing the claim that statistical knowledge suffices for legal liability when the fact that the defendant "knows" is a circumstance rather than a result. However, in an important recent "willful blindness" case, the Court of Appeals for the Ninth Circuit endorsed an individualized knowledge conception of the knowledge required for conviction of a federal drug offense, and rejected a statistical knowledge conception. *See* United States v. Heredia, 483 F.3d 913, 923-24 (9th Cir. 2007). The concurring opinion in the case expressed the worry that the majority's test would permit very broad criminal liability for knowingly transporting drugs, based only on (what I have called) statistical knowledge:

> [B]usinesses [such as FedEx] are not "knowingly" transporting drugs in any particular package, even though they know that in a volume business in all likelihood they sometimes must be. They forego inspection to save time, or money, or offense to customers, not to avoid criminal responsibility. But these reasons for not inspecting

A good way to begin answering the statistical knowledge conundrum is to revisit the three critical functions 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.[31] 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.[32] 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.[33] Moreover, in tort law, punitive damages will often be available in individualized knowledge cases, but they will not normally be available in statistical knowledge cases.[34] Statistical knowledge cases, if culpable enough to deserve legal sanction, will ordinarily be classified

---

are not the ones acceptable to the majority . . . . The majority opinion apparently makes these businesses felons despite the fact that Congress did not.

Id. at 928 (Kleinfeld, J., concurring). The majority opinion responded:

Equally misplaced is the concurrence's concern about FedEx and similar package carriers. The fact that a tiny percentage of the tens of thousands of packages FedEx transports every day may contain contraband hardly establishes a high probability that any particular package contains contraband. Of course, if a particular package leaks a white powder or gives any other particularized and unmistakable indication that it contains contraband, and the carrier fails to investigate, it may be held liable – and properly so.

Id. at 920 n.10 (majority opinion) (citation omitted).

[31] 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. Consider, for example, *Laidlow v. Hariton Machinery Co.*, 790 A.2d 884 (N.J. 2002). The court held that the "intentional" tort exclusion is broad enough to permit a worker to sue an employer who "knew that it was substantially certain that the removal of the safety guard would result eventually in injury to one of its employees," but the court required, in addition, proof of other aggravating contextual factors that suggest the legislature could not have intended to immunize such conduct from tort liability. *See id.* at 897-98.

[32] *See supra* note 24 and accompanying text.

[33] The point in the text also applies in tort law. Consider again the examples mentioned in the *Restatement (Third) of Torts*. *See supra* note 8. 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.

[34] Some famous tort cases – especially the Ford Pinto case – seem to be striking counterexamples to this proposition, but I will argue that they are not. *See infra* Part III.B.

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.

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 it will not sanction Agatha at all, even though her justification is arguably not much more compelling than Clara's.

Why are statistical knowledge and individualized knowledge treated so differently? And what, exactly, *is* the relevant difference? The answers to these questions are vitally important, not only in differentiating more culpable knowing actions from less culpable reckless or negligent ones, 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. In a *paradigm individualized knowledge case* such as Clara, the actor knows that there is a very high probability that she will cause serious harm to a single person. The decision to take an action that one realizes is very likely to cause serious harm to another is virtually always unjustifiable. Only a small number of extremely narrow justifications, most prominently self-defense and necessity, will suffice.[35] All else being equal, one who knowingly creates a high risk of harm needs a weightier justification to deserve exemption

---

[35] As noted above, many states that recognize a necessity defense do not permit it as a defense to murder. *See supra* note 28. 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. To be sure, there are independent 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, if a jurisdiction retains the limitation, and if statistical knowledge counts as individualized knowledge, then the limitation is especially objectionable; for then every corporate decision to market a product known to cause a small number of deaths in the population of users would seem to qualify as unjustified murder.

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 and contentious 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 an actor needs a 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 interest *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 paradigm individualized knowledge case, the law properly treats the actor harshly – more harshly than an actor whose conduct is similarly unjustified but who demonstrates 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.[36] Knowingly taking a high risk of seriously injuring someone for the purpose of one's career usually demonstrates more serious fault 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 culpability category, other than individualized knowledge, should apply? Recklessness? Negligence? Or no culpability at all? The short answer is that any of the above could be the appropriate culpability level, depending on the context. Let us see why.

Consider two stylized statistical knowledge cases:

Larry (frequent careful driver)

Larry routinely drives carefully, very near the speed limit. He drives 15,000 miles per year. He is unusually well-informed and knows the aggregate risk of injury that this poses over his lifetime of driving: 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.[37]

---

[36] 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. I do not, however, pursue the issue here.

[37] These are invented 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 Frequently Asked Questions*, NAT'L SAFETY COUNCIL, http://www.nsc.org/news_resources/

Maxwell (frequent speeder)

Maxwell routinely drives ten miles per hour *over* the speed limit. Like Larry, he drives 15,000 miles 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 thirty occasions, (b) minor personal injury to someone else on at least five occasions, and (c) moderate personal injury on at least one occasion.

Despite Larry's statistical knowledge that he will, over time, cause harm to others, he should not be considered as culpable as an actor with individualized knowledge. Indeed, he probably should not be found legally responsible at all.[38] Maxwell, on the other hand, is a more complex case. In an important sense, he is 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. Nevertheless, his repetition of an unjustifiable action legitimately bears on his culpability, as we will see.

These two examples suggest a preliminary answer to the question of how to categorize statistical knowledge. Statistical knowledge ordinarily belongs with recklessness or a lesser culpability level, not with individualized knowledge. This conclusion is also supported by our examination of 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*

Resources/res_stats_services/Pages/FrequentlyAskedQuestions.aspx#question10 (last visited Oct. 21, 2011). The lifetime risk of death in an automobile accident is 1 in 88. *See Lifetime Odds of Death for Selected Causes. United States, 2007*, NAT'L SAFETY COUNCIL, http://www.nsc.org/NSC%20Picture%20Library/News/web_graphics/Injury_Facts_37.pdf (last visited Oct. 21, 2011).

This Article focuses on the legal and moral significance of an actor's subjective awareness of a risk. But what about "objective" beliefs about risks? How should we treat an actor 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 Article. 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.

[38] 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. This issue, however, is outside the scope of this Article.

of the same harm but who knows that the risk will befall a particular individual. Compare the following scenarios. In the first, a delivery service decides to use automobiles rather than bicycles in a particular municipality, despite statistical knowledge that in the long run this will cause more injuries to pedestrians but will greatly speed delivery time. In the second, a motorist or bicyclist, on a single occasion, decides 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. We properly demand a heavier burden of justification when an actor knowingly creates a high, rather than a low, risk to a victim on a single occasion. It does not follow, however, 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.[39]

   To see the power and limits of these two arguments, let us turn to a more formalized explanation of why we should ordinarily treat statistical knowledge and individualized knowledge differently.

---

   [39] Or, to put the matter differently, how one frames the scope of the risk of an activity is quite arbitrary: one might describe the risk 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. 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. *See* Simons, *supra* note 16, at 1218-21.

   This point 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 16, at 1144-47.

   Or consider the infamous Ford Pinto case, *Grimshaw v. Ford Motor Company*, 174 Cal. Rptr., 348 (Cal. Ct. App. 1981), 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. *See* Schwartz, *supra* note 13, at 1015-19; Simons, *supra* note 22, at 485. 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. *See* Schwartz, *supra* note 13, at 1022. 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. *See* Schwartz, *supra* note 13, at 1020. That sounds like an insignificant percentage. We could, however, 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 deaths. *Id.*

B.   *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 merely 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 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, 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:

Invariant Culpability when Acts are Aggregated (ICAA):

Suppose: 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: 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 produce benefits will typically change; he will recognize that those harms 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 will be liable 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 over either time or space. If Alicia only drives once in her lifetime, while Bella

drives 10,000 times, and Carmen 300,000 times, ICAA provides that their degree of culpability for resulting harm should be considered the same as long as their driving behavior is the same.[40] And if a company expands its fleet of drivers 10-fold, 100-fold, or 10,000-fold, but otherwise maintains the same safety policies, we should not conclude that it is more culpable for the harms its drivers cause simply because the aggregate risks are much greater.[41]

It might seem that statistical knowledge should be treated like individualized knowledge at least in the special case where the activity in aggregate is likely to cause *one or more* instances of serious physical injury or death. This argument is flawed, however, 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.[42] Over a span of a second, few activities will. Yet neither description is more 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 thus far has analyzed how ICAA applies to a particular number of "drives," but the term obviously is inexact. To be more precise, we could try to individuate not distinct "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

---

[40] 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 Article, assume that the characteristics of a single "drive" are reasonably well-defined, so that we can meaningfully compare different actors, each of whom takes a single (relevantly similar) "drive," and can also compare an actor engaging in one "drive" to that actor engaging in many (relevantly similar) "drives."

[41] 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).

[42] 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, *Conundrum for the Day*, VOLOKH CONSPIRACY (Jan. 14, 2004, 9:04:05 AM), http://volokh.com/2004_01_11_volokh_archive.html#107409984581 699780 (arguing that in such a hypothetical, we would get the absurd result of "banning driving for people with sufficiently long lives"); Tyler Cowen, *Would Potential Immortals Be Risk-averse?*, MARGINAL REVOLUTION (Jan. 15, 2004, 8:04 AM), http:/www.mar ginalrevolution.com/marginalrevolution/2004/01/would_potential.html.

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, and steering farther away at the fourth.[43] 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 has engaged in one episode of risky behavior, while Arthur has engaged in three comparable episodes.[44] For purposes of our inquiry, this notion of comparable "episodes" is sufficiently precise, and so is the notion of comparable "drives," if suitably qualified.[45]

---

[43] A single drive might also endanger no persons, one person, or many persons. Indeed, physically similar acts within one "drive" could differ in the type of risks of harm that they create and in the number of persons exposed to those risks.

This 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 it 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, invariant culpability when risk-exposures are aggregated (ICREA). *See infra* note 89.

[44] 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 an infamously 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.

[45] 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 theory. *See, e.g.*, LARRY ALEXANDER & KIMBERLY KESSLER FERZAN WITH STEPHEN MORSE, CRIME AND CULPABILITY: A THEORY OF CRIMINAL LAW 244-60 (2009) [hereinafter ALEXANDER & FERZAN]; MICHAEL S. MOORE, ACT AND CRIME: THE PHILOSOPHY OF ACTION AND ITS IMPLICATIONS FOR CRIMINAL LAW 305-90 (1993) (analyzing identity conditions of actions and individuation in the context of double jeopardy).

Fortunately, for purposes of this Article, 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? ICAA gives the presumptive answer: no. I 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.

What deeper theory of permissible risk-creation and of punishment underlies ICAA? Perhaps the most straightforward defense of ICAA is provided by a consequentialist[46] 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, if we are employing the Learned Hand test,[47] and if we conclude that it is not negligent for Larry to drive at a moderate speed rather than walk, then we must have first determined that the burden of not driving at moderate speed on a given occasion is greater than the marginal costs, costs that prominently include the risks of injuring others. 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.[48]

One further reason justifies agnosticism about how to individuate "acts." The only concept that we really need 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 repeatable (and potentially culpable) conduct in which the actor creates risks of harm. The episode could, but need not, be an "act." Let me explain.

Individualized knowledge, as I define it, entails that an actor believes he is imposing a risk of harm on the particular victim greater than some threshold probability T (e.g., 75%). Thus, 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 he believes creates a risk less than T to an individual victim, while A2 engages in a longer episode that he also believes falls short of creating a risk as high as T to an individual victim, then (a) neither has acted with individualized knowledge, and (b) A2's culpability is presumptively no greater than A1's.

[46] This form of consequentialism, however, is truncated, for 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, *supra* note 16, at 1184-85.

[47] The Learned Hand test need not be understood in consequentialist terms, however. *See* Kenneth W. Simons, *The Hand Formula in the Draft* Restatement (Third) of Torts*: Encompassing Fairness as Well as Efficiency Values*, 54 VAND. L. REV. 901, 909-14 (2001); Simons, *supra* note 16, at 1221 ("[T]he [Learned Hand test] can accommodate both sensitive consequentialist and tough-minded deontological accounts of negligence, each of which is much more plausible than unqualified consequentialist or unqualified deontological accounts.").

[48] *See* RICHARD A. POSNER, ECONOMIC ANALYSIS OF LAW 261 (8th ed. 2011), for the following example:

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 operations – makes B high.

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 the following: it is only justifiable to impose a risk of serious injury on others so long as those others consent to the risk, benefit from the risk, or 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. For analogous reasons, most other plausible nonconsequentialist criteria are consistent with ICAA.[49]

Consider an abstract depiction of the argument:

The ratio of B to PL is unaffected by the scale of the potential injurer's operation, and it is the ratio that distinguishes intentional and unintentional torts in an economically relevant sense.

The analysis in the text requires qualification, however. It is possible that a 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 *is* worth regulating if the marginal costs of regulation decline sufficiently with the increased number of acts. For example, the State might not find it worthwhile to prosecute a petty theft unless the offender is known to have committed numerous small thefts.

[49] 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 Simons, *supra* note 16, at 1193-1202.

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 highly problematic. Some nonconsequentialist legal academics argue that it is wrongful to impose "substantial" or "significant" risks. *See* Keating, *supra* note 2, at 180; Richard W. Wright, *Right, Justice and Tort Law*, in Philosophical Foundations of Tort Law 159 (David G. Owen ed., 1995). Such arguments have to be formulated very carefully lest they trigger the "arbitrary framing" objection that I have noted.

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; the numbers represent magnitudes of values that cannot really be quantified with any precision.[50] "20,000" is simply an abstract measure of the disvalue of creating a 70% risk of causing personal injury. The example does not assume a consequentialist justification for the legal sanction or regulation of the risky conduct, though it indeed assumes that the justification, whether consequentialist or nonconsequentialist, requires some type of balancing of socially recognized risk and socially recognized benefit. Despite its abstraction, the chart remains highly instructive, as we will see.

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 he knows there is a 70% probability that over 20,000 drives his driving will cause personal injury to someone, he should not be treated the same as

---

[50] Moreover, whether one is employing economic analysis or another approach (including nonconsequentialist principles) to assess the justifiability of risk-creation, one should focus on the *marginal* risks and benefits of a particular level of care. *See* Simons, *supra* note 16, at 1172. For simplicity, I do not emphasize this marginal feature in the charts, but it would be possible to modify the charts to reflect it.

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 79 (2010).

those actors possessing individualized knowledge. His culpability does not increase simply because he drives more frequently.

ICAA sometimes needs to be qualified 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 farm or manufacturer could justifiably produce a small amount of pollution in light of the social benefits of its activity does not mean that it can justifiably produce a much larger amount of pollution, even if the social benefits are also proportionately greater.[51] 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. These types of cases are exceptional, however, 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 actor is free to impose a risk of harm. Aggregate property damage and physical injury are social harms that normally increase in rough proportion to increases in such risky activities as 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 repetition preserve whatever level of fault the actor displays in a single action? Or can repetition *increase* the level of fault?[52] Recall the example of Maxwell, who

---

[51] For discussion, see 1 JOEL FEINBERG, THE MORAL LIMITS OF THE CRIMINAL LAW: HARM TO OTHERS 225-32 (1984).

A range of scenarios needs to be addressed in order to analyze fully 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.

[52] 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 remains negligent or reckless, whether it applies the policy to 10, 1,000, 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 infra* text

routinely chooses to drive ten miles per hour over the speed limit. In any individual drive, he is negligent and at worst reckless. 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? At 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 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.[53]

Yet this conclusion, that ICAA should extend to the benefit of faulty actors, is 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 repeatedly takes *unjustifiable* risks or otherwise demonstrates culpability. Here, it seems, the whole can be more than the sum of its parts. On this view, if Maxwell speeds twenty times, his conduct displays greater aggregate fault than if twenty different, unrelated individuals each speeds once.

Consider the fact that the criminal laws of most states punish recidivists especially harshly.[54] 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 of his offenses will exceed – often by an enormous amount – the sum of the distinct punishments that would otherwise be warranted if each of his offenses had been committed by different

accompanying notes 88-89.

[53] 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.

[54] 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 below.

individuals.[55]   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 these dangerous acts would only warrant compensatory damages had separate individuals committed each act.[56]

A consequentialist has good reason not to apply ICAA to faulty conduct. If Maxwell repeatedly acts unjustifiably, a greater than usual sanction will often be needed to achieve specific deterrence, since he by hypothesis has not been deterred by the normal threatened sanction for his prior acts. 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. For Larry, however, who repeatedly acts *justifiably*, consequentialism requires adherence to ICAA. As we have seen, a cost-benefit analysis suggests no reason for re-characterizing his culpability as faulty rather than faultless simply because he repeats his acts.

Does a *nonconsequentialist* also have good reason not to apply ICAA to wrongful and unjustifiable conduct? The answer is less straightforward. One important complication here flows from an empirical observation: when someone repeats faulty 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. On closer inspection, however, such cases are not really a counterexample to ICAA; they

---

[55] But 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 would 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, 215 (2005); *see also* ALEXANDER & FERZAN, *supra* note 45, at 254-57. 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.

Many complicating factors are at play here beyond principles of just deserts. For example, a judge might give the one-time speeder a more serious 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.

[56] 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. 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 infra* text accompanying note 62.

are simply instances in which, as an empirical matter, ICAA's *ceteris paribus* condition does not apply. Let me explain.

As discussed 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 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.[57] Although committing a dangerous act once may only be negligent, repeating the act will often be reckless, simply because the repeating actor has become cognizant of the risks.[58] Similarly, when the conduct repeats, a legal decision-maker will have more *confidence* in his judgment 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.[59] And since he has more opportunities for reconsideration of his acts and more opportunities to take mitigating precautions, his failure to take such steps

---

[57] 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 makes him less culpable. Yet 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).

[58] 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 context – for example, that the brakes of a rental car might be defective. But if he rents cars every month and on one occasion discovers this problem, perhaps he has a duty to check the brakes thereafter.

[59] In the case of an impulsive, sudden reaction, 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* ROBERT NOZICK, PHILOSOPHICAL EXPLANATIONS 392 (1981) ("[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* Michael S. Moore & Heidi M. Hurd, *Punishing the Awkward, the Stupid, the Weak, and the Selfish: The Culpability of Negligence*, 5 CRIM. L. & PHIL. 147, 153-56, 192-93 (2011). 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.

suggests a more deeply considered elevation of his own interests over the welfare of others and thus adds to his culpability.[60]

Nevertheless, this empirical gloss is not the entire nonconsequentialist story. Apart from the truth or falsity of these empirical contingencies, both the recidivist premium and similar policies of aggravation are defensible in principle. Specifically, it seems justifiable to

> (1) punish recidivists more harshly than they would deserve if they had only committed their most recent crime;

> (2) impose especially severe punishment on actors who successively create multiple risks of harm in a single extended criminal episode;[61] and

> (3) impose punitive damages on corporations that continue to market a dangerous product even after they realize the danger.

The plausibility of these positions suggests that ICAA is *normatively* unpersuasive in this context, even though it might otherwise apply. An actor who persists in wrongdoing is legitimately viewed as committing a wrong more serious than the composite wrong that is constituted by the sum of the individual unjustifiable acts. Why is this so?

When an actor repeatedly commits unjustified acts, he often reveals the culpability of *defiance*. This defiance of legal and social norms is a distinctive type of wrong that transcends the culpability of the individual acts and might properly warrant enhanced legal sanction.[62] The whole can indeed be greater than the sum of the parts.

The question remains: just how culpable are such "repeat offenders"? 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

---

[60] 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.

[61] *See, e.g.*, People v. Gomez, 478 N.E.2d 759, 761-62 (N.Y. 1985) (affirming conviction of "depraved mind murder" because of the defendant's willingness to expose numerous victims to lethal harm by his highly dangerous driving).

[62] I do not have space here for a full discussion of the nature of defiance as a distinct form of culpability. One valuable discussion is Jean Hampton, *Mens Rea*, 7 SOC. PHIL. & POL'Y 1, 15 (1990) ("[A] culpable agent is one who chooses to defy what she knows to be an authoritative moral command in the name of the satisfaction of [her] wishes, whose satisfaction the command forbids."). 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 P. GREEN, LYING, CHEATING, AND STEALING: A MORAL THEORY OF WHITE-COLLAR CRIME, 114-26 (2006) (discussing relevance of disobedience to the law in assessing the moral content of certain crimes). *But see* Youngjae Lee, *Recidivism as Omission: A Relational Account*, 87 TEX. L. REV. 571, 592-608 (2009) (rejecting "defiance of authority" as justification for punishing recidivists).

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 this variation to Ned taking a single highly dangerous 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.[63]

In the Maxwell/Ned scenarios, acting with statistical knowledge, although unjustifiable, is much less unjustifiable than acting with individualized knowledge. The individualized knowledge case is characterized by a much greater disproportion between risk and conceivably justifying benefit. Assume that the actor's reason for taking the risk is the same in both cases – to arrive home earlier in order to relieve a babysitter and care for a sick child.[64] In highly stylized form, the two cases and the case of Larry, the faultless driver, can be analyzed as follows:

---

[63] One significant complication for this analysis is whether we reject or accept moral luck. Moral luck is the idea that the occurrence or nonoccurrence of a harmful result should matter to culpability. If we reject moral luck, then Maxwell might be punished equally for all his acts of endangerment, whether or not they cause injury, 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.

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

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

[64] Even if the "benefit" is of little or no social value, the structure of the analysis does not significantly change, though the "benefit" numbers would be much lower. (Suppose the driver desires to catch the beginning of a televised basketball game or seeks the pleasure of racing a friend to his house.)

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

As the chart reveals, Maxwell's conduct over 4,000 drives produces the same aggregate disvalue as Ned's conduct in a single drive, but a *much* larger social benefit. This dramatic difference legitimately bears on the appropriate legal sanction for each actor. Maxwell's 4,000 drives are, in this important way, much less culpable than Ned's single drive. Of course, in another way – the defiance exhibited by the repeated faulty behavior – Maxwell's 4,000 drives are more culpable.

D. *The Special Stringency Principle 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. Because of SSP, 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.[65]

Two more characters illustrate the point:

Alfa and Benna

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.[66]

How should we treat these examples?

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.[67]

Consequentialists are likely to treat Alfa-situations and Benna-situations the same, because each scenario produces the same expected net saving of lives – in the sense that each produces an identical 5:1 ratio of [expected lives saved: expected lives lost]. But many nonconsequentialists will treat the two scenarios differently in the way suggested in the quoted passage.[68]

We can represent SSP in graphic form:

---

[65] See Kenneth W. Simons, *Negligence*, 16 SOC. PHIL. & POL'Y 52, 65 (1999), from which the following discussion is taken.

[66] *Id.* I owe this example to Leo Katz.

[67] *Id.* Assume, in this example, that there is no point in time when Benna both realizes that she is *almost certain* to kill the pedestrian and is able to avoid doing so. The assumption is critical in order to distinguish Benna's conduct and culpability from Alfa's. Thus, imagine that Benna is driving on a dark street at night and will not be able to see any pedestrian until after her car collides with him, and she believes that by driving in this manner to the hospital, she is creating a 20% risk of death.

[68] Nonconsequentialists might also consider the risk imposition more justifiable as the absolute level of the benefits and the risks decreases further, even though 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** (of saving life)
Alfa, one drive [probably at fault]	1,000,000 [100% risk of causing one death]	5,000,000 [100% chance of saving *five* lives]
Benna, one drive [not at fault?]	200,000 [20% risk of causing one death]	1,000,000 [100% chance of saving *one* life]

Again, the example is stylized,[69] but it illustrates the basic point: under a nonconsequentialist approach, even if Benna's conduct is justified, more than a proportionate increase in the expected benefit is required to justify Alfa's conduct. 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 crucial variation: suppose Benna engages in *repeated* acts with a cumulative risk as great as Alfa's single act? Examine the last row in the chart below, a scenario in which Benna takes five drives:

---

[69] The chart assumes that the benefit of saving a life is precisely equal in value or moral urgency 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; 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.

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. 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 basic analysis.

Chart 4: Alfa, Benna, Repeat Benna

	Disvalue of risk (of causing death)	**Value of benefit** (of saving life)
Alfa, one drive [probably at fault]	1,000,000 [100% risk of causing one death]	5,000,000 [100% chance of saving *five* 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, five drives [not at fault?]	1,000,000 [almost 100% aggregate risk of causing one death]	5,000,000 [100% chance of saving *five* lives]

Is the culpability of "repeat Benna" (R-Benna), who takes five equally risky drives, closer to that of Alfa or closer to that of "one-drive Benna"?[70] The discussion earlier provides an answer. If "one-drive Benna" acts permissibly, then ICAA provides that 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 repeated 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."[71]

---

[70] A further complication is how to compute 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 probabilities. So the aggregate risk from five drives will be less than 100%. Specifically, the risk that Benna will cause death after five separate drives is $1-(1-.2)^5 = 1-(.8)^5 = 0.67$, or 67%.

This complication is ultimately not that important: 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's behavior 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!)

[71] It is also possible that as R-Benna increases the number of her drives, she develops a clearer awareness of the risks that she is running on any given drive, or in some other way her later drives are no longer relevantly identical to her first. *See* discussion *supra* notes 57-

Here is the exception. Suppose 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* drives past the child enough times to pose an aggregate risk to that child equal to the risk that Alfa poses 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*, like Alfa, is now imposing *all* of the risk of her activity on a single person. And that high concentration of risk is precisely why *Alfa's* conduct is especially difficult to justify.[72]

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. Under SSP, 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.[73] SSP nevertheless is consistent with ICAA, because ICAA simply does not apply when acts significantly differ in how they distribute the risk of harm. If R-Benna's five drives endanger different 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.

As another possible illustration of SSP, recall Eleanor, the builder who decided to forge ahead with her tunnel construction project despite her awareness that continuing the project requires lowering a fitting into place that will very likely kill a worker trapped below. She would properly be convicted

---

60.

[72] For a similar example, see ALEXANDER & FERZAN, *supra* note 45, 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.

[73] For another illustration, recall Maxwell and Ned, *supra* Chart 2. In that 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. However, if one can imagine a world in which Maxwell's thousands of drives all endanger only a single victim (for example, 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.

of murder, I believe, for knowingly and unjustifiably causing that death. But why should her punishment be so harsh? Is it not perfectly acceptable for a tunnel builder to proceed with great care to build a tunnel even if she predicts that this will cause at least one death? (Remember Agatha.) And is it not at worst negligent for a second tunnel builder to proceed, with less care than Agatha, to build a tunnel even if she predicts that she will create a somewhat greater risk of injury and death than Agatha? (Remember Bertha.)

Again, SSP helps explain the pattern. Eleanor has chosen a course of action that she knows is virtually certain to kill the trapped worker, and that choice requires an extraordinarily compelling justification. The other tunnel-building 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.[74]

But this deontological principle is a bit puzzling. Why should it matter that the risk is concentrated in this way? One possible explanation is that highly concentrated risk-exposures are especially likely to be nonreciprocal, disproportionately endangering the victim while disproportionately benefiting others; they are therefore especially likely to be unjust. Another explanation is

---

[74] The Keating/Scanlon "World Cup" example, *see supra* note 10 and accompanying text, should be analyzed similarly to the Eleanor example: 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.

Unlike the Alfa/Benna distinction however, the Eleanor/Agatha distinction rests on an additional factor – the difference in the causal path by which Eleanor and Agatha bring about the harm. Notice that in the Eleanor example, it would be impermissible and highly culpable for the builder to lower the pipe even if this act endangered a dozen workers, one of whom was likely to die. I do not analyze this additional causal factor here.

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: RIGHTS, RESPONSIBILITIES, AND PERMISSIBLE HARM 29-30 (2007) [hereinafter KAMM, INTRICATE ETHICS]; 2 F.M. KAMM, MORALITY, MORTALITY 303-06 (1996) [hereinafter KAMM, MORALITY, MORTALITY]. Deontological principles, she plausibly argues, forbid both the policy and the installation of the device. *See* KAMM, INTRICATE ETHICS, *supra*, at 30 ("In general, we cannot permissibly bargain away our moral status not to be treated in certain ways in order to increase our life prospects nor to minimize overall rights violations."); KAMM, MORALITY, MORTALITY, *supra*, 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.

that the potential victims might agree ex ante to accept lesser risks but not highly concentrated risks.

These explanations help, but they also prompt three counterarguments. 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%?[75] Second, this approach seems to count some lives as more worthy than others. Why is the life of Eleanor's victim more valuable than the life of Agatha's or Bertha's victim? Third, the approach seems to ignore the preferences of those affected. After all, the workers might all prefer, ex ante, 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.

I believe that SSP can withstand these counterarguments. Concerning the first and second points, the question we are addressing is how culpable the actor is in various legal 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 possessing only statistical knowledge (e.g, Agatha or Bertha) lacks that type of foresight. An actor such as Eleanor, however, is indeed highly culpable when she chooses to sacrifice a *known* victim for the sake of more diffuse social benefits.[76] And even though the lives

---

[75] *See* ALEXANDER & FERZAN, *supra* note 45, at 65 (discussing "Sam"); Richard Brook, *Statistical and Identifiable Deaths*, *in* PHILOSOPHY AND ITS PUBLIC ROLE 167, 177 (William Aiken & John Haldane eds., 2004) (asserting that a builder such as Agatha or 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 SSP. 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? *See 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 that it is implausible to make a project's permissibility depend on whether the victims of the project are known. *See 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 it is more culpable to knowingly impose the first risk than the second. But I concede that these indirect knowledge (via third parties) examples are difficult to sort out. Compare a case in which the builder knows that a manager on the project will knowingly kill one worker (in the manner, say, of Eleanor), and 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.

[76] 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; for example, SSP could be restricted to situations in which the actor subjectively believes that he is imposing a highly concentrated

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 – namely, what is the culpability of the actor who causes such a victim's death?

One way to appreciate the difference in perspectives is to posit a different hypothetical: 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. Suppose by the same budgetary investment 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 (2) prevent *one* impatient driver such as Clara from knowingly killing a victim. And suppose 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 or Maxwell.[77]

The third counterargument, that a deontological constraint such as SSP sometimes requires that we not honor unanimous ex ante preferences, might seem to be a conclusive objection against SSP. But it is not. A distinctive and attractive feature of nonconsequentialist views is that they sometimes disallow an option that maximizes aggregate welfare.[78] What justifies this seemingly paradoxical element of nonconsequentialism? Both arguments of principle and widespread intuitive responses to examples should be consulted here.[79] Of course, consequentialists will reject this entire approach. Indeed, if they focus only on the bottom-line "body count,"[80] they will consider Eleanor's conduct,

---

risk *and* a reasonable person (or some other idealized actor) would also believe that the risk is highly concentrated. I do not pursue these complications.

[77] *See* Leo Katz, *Incommensurable Choices and the Problem of Moral Ignorance*, 146 U. PA. L. REV. 1465, 1475 (1998) ("[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.").

[78] A number of philosophers and legal scholars have emphasized the point. See KAMM, MORALITY, MORTALITY, *supra* note 74, at 290-310; T.M. SCANLON, WHAT WE OWE TO EACH OTHER 82-87 (1998); Katz, *supra* note 77, at 1474-75. Thus, in the famous Transplant Case, nonconsequentialists would consider it unjustifiable to take bodily organs forcibly from one person, thus knowingly causing his death, in order to save the lives of five. Carol S. Steiker, *No, Capital Punishment Is Not Morally Required: Deterrence, Deontology, and the Death Penalty*, 58 STAN. L. REV. 751, 761-62 (2005). Indeed, it is also unjustifiable merely to take a kidney forcibly from A (creating a health risk to him but not killing him) in order to save the life of B. *See* Charles Fried, *The Value of Life*, 82 HARV. L. REV. 1145, 1430-31 (1969). 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.

[79] Frances Kamm famously relies heavily on intuitive responses to hypotheticals in justifying her "intricate ethics." *See* KAMM, INTRICATE ETHICS, *supra* note 74, at 3-8; KAMM, MORALITY, MORTALITY, *supra* note 74, at 9-12.

[80] I say *if* they focus only on the bottom line because consequentialists can incorporate

as well as the conduct of our speeding-to-the-hospital drivers Alfa and Benna, permissible.[81] But most people judge this type of unqualified, relentless cost-benefit calculus as morally objectionable, and our legal system reflects that judgment.

E.   *How Should We Distinguish Statistical from Individualized Knowledge?*

Up to this point, most of the examples have been easy to sort into either the statistical knowledge or the individualized knowledge category. Many other examples, however, are difficult to categorize. And even when we can confidently classify an example, the implicit criterion that drives our intuition is sometimes obscure. For example, some cases that intuitively seem very similar to paradigm cases of individualized intention or knowledge rather than mere statistical knowledge nevertheless involve an extended class of potential victims or an extended time dimension. An instance is the Time Bomb example. 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. A similar pair of examples can be constructed for harms that the actor knowingly rather than purposely causes.[82] But how do such cases differ from genuine statistical knowledge cases?

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

Consider two recent efforts to develop workable criteria. Neither entirely succeeds, but both illuminate the problem. Additionally, the first effort

---

other factors into 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 they rely on the contingent preferences of the public, yet the public sometimes unjustifiably rejects valid deontological principles. Moreover, closer alignment can be achieved only if the policy is known to the populace; yet, by deliberately hiding a morally questionable policy from public view, government might confer *greater* social utility on the policy. It seems illegitimate to defend a policy whose social value depends on such secrecy.

[81] For example, Kaplow and Shavell strongly imply that they would find Eleanor's action morally and legally permissible. *See* Kaplow & Shavell, *supra* note 4, at 1277 n.757.

[82] For example, 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. 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?

identifies an important 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.[83] Tort doctrines requiring "intent" often treat the term as encompassing both purpose and knowledge. But what counts as "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.[84] They suggest that the basic distinction is between

(1) the harm predicted to result from an individual act;

and either

(2a) the harm predicted to result from *a course of repetitive conduct*; or

(2b) the harm predicted to result when an individual act that creates a continuing danger *interacts 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.[85]

An example of category (2a) is a scenario involving a baseball player who hits 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 during the season.[86]

Henderson and Twerski persuasively treat category (2a) as instantiating mere statistical knowledge, not genuine individualized knowledge. This category is usually a straightforward application of the ICAA principle.[87] When a baseball player hits a single fly ball, he will at most be aware of only a

---

[83] *See* Henderson & Twerski, *supra* note 2, at 1141-43.

[84] I say "essentially" agree because they frame the question as whether "the consequence complained of must result directly and proximately (both temporally 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* Anthony J. Sebok, *Purpose, Belief, and Recklessness: Pruning the* Restatement (Third)*'s Definition of Intent*, 54 VAND. L. REV. 1165, 1169-72 (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. 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").

[85] *See* Henderson & Twerski, *supra* note 2, at 1143.

[86] *Id.* at 1141-42.

[87] I hedge my endorsement of their category (2a), however, because they do not include the qualification I endorse: the repetitive conduct must consist of an aggregation of acts that do not differ in any relevant respect from the single act.

small risk that the ball will injure a fan. Repetition of such acts does not elevate his culpability to the level of individualized knowledge.

But the authors' other 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, their elucidation of that category is too broad, treating too many cases as cases of statistical knowledge when some should be treated as cases of individualized knowledge.

1. The Principle of ICREA (Invariant Culpability when Risk-Exposures are Aggregated)

Here is an example that fits comfortably within category (2b):

Jungle Gym

A playground equipment company designs and 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 injures? Certainly not. The mere fact that repetitive use of the structure by thousands of children is likely to lead to injury does not mean that the company acted tortiously, and it surely does not mean that the company is an *intentional* tortfeasor. Compare a company that manufactures an "extreme sport" version of the jungle gym for elementary schools, with slippery surfaces and a twenty-foot drop to the ground. In this 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 whether the manufacturer's design of the equipment balanced risk and benefit permissibly, negligently, or quite recklessly, we should render the same judgment even when the instances of use multiply.[88]

---

[88] If, however, a jungle gym will be heavily used only by one or two actors, arguably the highly concentrated risk would trigger SSP. 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

However, Jungle Gym is also *different* from the ICAA cases we have discussed thus far because the repetition of risk exposures and accidents in Jungle Gym is not due to a repetition of dangerous acts by the company making the equipment; it is due to the fact that, over time, numerous users interact with the potentially dangerous equipment, thereby increasing the aggregate amount of risk exposure. This scenario illustrates a second "invariant culpability" principle: "Invariant Culpability when Risk-Exposures are Aggregated" (ICREA).[89] ICAA instructs us that the culpability properly attributed to a single act should also be attributed to multiple acts of precisely the same type, while ICREA instructs that the culpability properly attributed to a single instance 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 that Henderson and Twerski offer:

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."[90]

---

if the risk on any single occasion of use was very high (for example, if the user could easily increase the height of a climbable structure on the jungle gym to twenty feet or more).

[89] We can spell out ICREA in a parallel fashion to the definition of ICAA (with the differences noted in *italics*):

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

Suppose: 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: 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 category (2b) in one significant respect: (2b) seems to require the interaction of the original act with the repetitive *acts* of other persons, while ICREA only requires that multiple persons be exposed to risks of harm. The broader principle is more sensible, for it properly encompasses passive victims who do not "act." For example, suppose Fanny sets off fireworks for the benefit of a passive crowd of spectators. ICREA instructs that the permissibility or culpability of her conduct is the same, whether 50 or 5,000 people are enjoying the spectacle.

[90] Henderson & Twerski, *supra* note 2, at 1142.

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; the injury should instead be analyzed under standard negligence principles.[91]

I disagree with the authors' treatment of Removed Safety Guard as just a straightforward example of category (2b) or ICREA and thus merely an instance of statistical knowledge. The risk in the example is concentrated on a relatively small group of workers, many of whom will be repeatedly exposed to that risk. As we have seen, ICAA 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 ICREA. Both principles are qualified by SSP. At the extreme, imagine that only one worker operates the machine: if, over time, the lack of a guard creates a virtual certainty that the worker will be harmed, his plight is almost as dire as that of a worker who, the employer can predict, will be harmed the very first time he works with the unguarded machine.[92]

In short, Henderson's and Twerski's approach is illuminating, for it reveals that foresight of likely future harm resulting from one's conduct can sometimes be characterized as mere statistical knowledge even when the conduct consists of a single act, rather than a series of acts. ICREA must supplement ICAA. In endorsing an extremely broad version of ICREA, however, their approach is unpersuasive.

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

---

[91] *Id.* at 1142-43. Henderson & Twerski explain, by way of summary, "[F]or 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.

[92] The example is problematic for two additional reasons. First, the extended time period 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. The example might be, in short, a potent illustration of the culpability of defiance. *See* discussion *supra* note 62.

Second, 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 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 – triggering liability, for example, or no longer permitting defenses of consent or assumption of risk. *See* discussion *supra* note 51.

> 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 becomes 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.[93]

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 ICAA, ICREA, and SSP.

2.   Statistical Knowledge of a Risk That Is Extremely Unjustifiable

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. That is, 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 the harm will occur. Recall the Time Bomb example. It seems

---

[93] RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL OR EMOTIONAL HARM § 1 cmt. e (2010). Comment e contains two other illuminating observations. If a company knows that its products will, because of the inevitable limits of quality control, contain some manufacturing defects that will cause physical injury, this might support strict liability but will not support intentional tort liability. *Id.* And in other situations, knowledge even of relatively certain harms is consistent with the absence of any tort liability. *Id.* The comment thus implicitly endorses ICAA.

The Reporter's Note to this comment shares my criticism of Henderson's 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 tort tend to find that the employer's knowledge in such a case satisfies the substantial-certainty standard.

RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL OR EMOTIONAL HARM § 1 cmt. e, reporter's note at 16 (2005); *see also* Laidlow v. Hariton Mach. Co., 790 A.2d 884, 898 (N.J. 2002).

absurd to deny that the terrorist knowingly wounds or kills and is thus 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.[94] 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.

One way to explain away this apparent counterexample is that the terrorist most likely *intends* to kill or wound. And we ordinarily 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 cause harm or one does not. Thus, there is no "statistical purpose" problem analogous to the statistical knowledge problem.

But the example is not so readily dismissed, for we can easily fine-tune it so that it is clear that the terrorist acts with knowledge and not purpose. 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, but he is quite aware that the means he has chosen is very likely to cause serious injury or death to one or more people. So modified, the example remains troubling, for it is not easily distinguishable from a standard individualized knowledge case, such as a driver knowingly running down a pedestrian. To be sure, the victims of the bombing might be difficult or impossible for the actor to identify in advance. But ex ante anonymity is of questionable legal or moral relevance,[95] and in any case it 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 months or years after a car crash.[96]

---

[94] *See* Heinzerling, *Knowing Killing*, *supra* note 11, at 530-31 (providing an analytically similar example of an actor who randomly distributes a lethal but undetectable poison in a small percentage of a large population).

[95] *See infra* APPENDIX.B.

[96] Nor does a significant physical distance between the actor and the instrumentality causing the harm plausibly distinguish the cases. The terrorist might detonate the bomb by means of a remote control from another room or even from another city. But in a standard case, too, the act might predictably have some physically distant effects. Suppose a driver collides with a truck transporting chemicals and thus causes widespread environmental harm.

If analyzed honestly, however, 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%, 70%, or 95%), in Time Bomb the actor might well believe that the risk to each individual potential victim is much less than T.[97]

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.

Similarly, 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. His culpability is almost as great as in a scenario where he removed the guard while employing only one worker, knowing that the worker was highly likely to suffer injury. Again, it is plausible 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* suggests, as we have seen.[98]

In short, Time Bomb and Removed Safety Guard show that it is entirely defensible for the law to endorse a capacious definition of individualized knowledge, embracing not only paradigm cases where the actor recognizes a probability of greater than T that an individual victim will be harmed but also cases where the actor is aware of a probability greater than T that one of *a small numbers of individuals* will be harmed. Under this approach, no bright line distinguishes individualized from statistical knowledge. And for similar reasons, SSP is probably best understood, at least for purposes of legal sanctions, as imposing a duty that becomes *continuously* more stringent as the concentrated risk that the actor believes he is creating increases, rather than a special duty that is triggered only when a defined *threshold* risk of harm is reached.[99]

---

[97] For another example, compare the following situations:
- RR1, who plays involuntary Russian Roulette with a single victim and fires multiple shots, which in aggregate impose a risk greater than T, and
- RR2, who plays Russian Roulette with multiple victims and fires multiple shots, which in aggregate impose a risk greater than T, but fires only one shot at each victim, each imposing a risk less than T.

RR1 acts with individual knowledge; RR2 does not.

[98] *See supra* note 93.

[99] Nevertheless, even if SSP is a continuous function, it is not a simple proportionate

The cases we just examined 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, which 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 that Jody Armour offers to show that a criterion of intention (and, by implication, a criterion of individualized knowledge[100]) can legitimately apply even over an extended period of time:

Programmed Rifle

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."

---

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 the act's 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, the duty 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, perhaps acts 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.

Does this analysis suggest that SSP is unnecessary? Perhaps, as Alexander and Ferzan would argue, 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. *See* ALEXANDER & FERZAN, *supra* note 45, at 67-68 (arguing that culpability is best identified as insufficient concern, which entails criminals "choosing to take risks to others' legally protected interests for insufficient reasons"). Although superficially attractive, this strategy is not an improvement. First, simply indicating that a range of relevant factors is relevant and requiring them to be balanced in a sliding scale provides an opaque formula. *See* Kenneth W. Simons, *Retributivism Refined – or Run Amok?*, 77 U. CHI. L. REV. 551, 581-83 (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.

[100] 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. *See* Armour, *supra* note 2, at 1160-61. 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.

> The rifle is connected to an automatic timer and is programmed to shoot once per day or week. [101]

Notwithstanding the significant time gap between the discrete acts, Armour claims, "[O]ur intuitions would demand that he be responsible for an intentional injury when the 'live' round is finally discharged into the crowd."[102]

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 Armour's example initially seem important – the actor's precommitment to a course of activity that will cause serious harm and the *extremely* unjustifiable nature of the conduct. In the end, it is the second feature that is critical in formulating a defensible criterion of individualized rather than statistical knowledge.

Consider first the actor's precommitment to a course of action. One can distinguish precommitment of different types and degrees,[103] but let us begin with the clearest case: an 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 injure many of its users. On reflection, precommitment cannot be a decisive factor. Businesses often 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 this should certainly not suffice to characterize

---

[101] The example is adapted from Armour's example. *See id.* at 1160-61. Armour offers it in criticism of the criterion, discussed above, proposed by Henderson and Twerski. *See supra* notes 83-86 and accompanying text.

[102] *Id.* at 1161.

[103] In the strongest form of precommitment, once a person has performed the initiating act, it is impossible for him to prevent the risk from materializing in harm. 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 still 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. Customary patterns of conduct could also be described as a weak form of "precommitment": the actor habitually drives to work rather than walks, bicycles, or takes a bus. But these distinctions are not critical to the main argument in the text.

them as knowing wrongdoers. It would hardly be proper to treat a large business as a knowing wrongdoer because it requires its sizeable sales force to drive – to drive responsibly but indeed to drive.

This is not to say that precommitment is never relevant to legal liability. If Domino's Pizza had a policy to encourage its delivery drivers to speed if necessary to deliver pizzas to customers within thirty minutes,[104] then the company's knowledge that this policy would lead to a predictable number of accidents and injuries is highly relevant to its liability. If that policy was sufficiently unjustifiable and if Domino's knew of its effects, punitive damages would be appropriate.[105] But that conclusion would be true whether drivers followed the policy once or a thousand times, assuming there is no difference in the company's knowledge of the risks or in other relevant factors.[106] In short, ICAA applies: 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 placing such cases within the category of individualized knowledge rather than statistical knowledge. If an actor either maliciously intends serious injury or knows that serious injury will occur, and if he is moved to act by a sadistic or other highly immoral reason, he is clearly among the most culpable wrongdoers. This is true even if the actor cannot know, when he precommits, the precise time or precise victim of the future injury. Indeed, one could extend the spatial and temporal dimensions of the example indefinitely, and it would still be proper 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 sometime in the next hundred years.

---

[104] *Cf.* Wauchop v. Domino's Pizza, 832 F. Supp. 1572, 1573 (N.D. Ind. 1993) (examining a policy of Domino's Pizza guaranteeing pizza delivery within thirty minutes that allegedly contributed to a death when a delivery driver struck the victim's minivan; the policy, however, explicitly required drivers to comply with traffic laws).

[105] 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 bear on the company's legal responsibility under *respondeat superior*.

[106] In practice, however, when an actor repeats his conduct, he will often have a more confident and certain belief about its harmful results. 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 supra* Part II.C. 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. *See supra* discussion note 62.

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 his immoral motivation warrant treating him as extraordinarily culpable. He is just as culpable as one who poisons a local water supply expecting a hundred people to die within the hour as a consequence.[107]

What is the deeper explanation for this "extremely unjustified" qualification? The answer draws upon the basic rationale for presumptively treating agents as serious wrongdoers when they act with individualized knowledge that they will cause serious harm. 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 in the crosswalk – 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 by getting him home earlier. Compare an analogous *statistical* knowledge case. 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 latter example so that the chemical *will* produce significant 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 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 sometime in the next hundred years.

---

[107] This example involves a single culpable act that exposes numerous people to risk. Thus, it exemplifies ICREA, not ICAA.

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

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 purposeful or knowing wrongdoers, such as the terrorist releasing the atmospheric poison.

Yet another way to appreciate why the degree of unjustifiability is so important 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 believes that the risk of harm is less than 70%.[108] There will, of course, be close cases. But it is sensible for the law to draw a relatively bright line: only a belief that one is creating a *high* degree of risk will trigger the distinctive doctrinal consequences of knowingly or "intentionally" causing injurious harm.[109] Murder by knowingly causing death is justifiably punished more harshly than reckless manslaughter; "intentional" battery, in the form of knowingly causing harm, is justifiably treated as a more serious tort than recklessly or negligently causing harm. Although borderline differences in the perceived risk of harm (e.g., 75% versus 65%) do not clearly warrant differential consequences of this sort, more substantial differences (e.g., 75% versus 10%) do.

Why are these differential consequences justified? 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

---

[108] In the text, I assume that "reckless" is understood in the purely cognitive sense, 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 supra* note 23 and accompanying text. This common use of a complex rather than purely cognitive conception of recklessness expands the legal difference between knowledge and recklessness. Thus, if D1 knowingly causes a serious harm, typically D1 is liable unless he can invoke an affirmative defense such as self-defense or necessity. But if D2 recklessly causes a serious harm, the burden of proof or production is different: D2 is liable for recklessness only if *the State or the plaintiff* can also show that D2's act was clearly unjustified.

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." The Model Penal Code seems to require that the actor believe that the risk is substantial, though the matter is not entirely clear. *See* Simons, *supra* note 21, at 189. But whether or not a jurisdiction narrows recklessness in this way, the basic analysis in the text comparing knowledge and recklessness remains valid.

[109] For a discussion of these consequences, see *supra* notes 24-28 and accompanying text.

have been willing to knowingly create the greater risk in pursuit of the same ends. And third, all else being equal, one who knowingly creates a higher risk needs a weightier justification 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 earlier discussion confirms. 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. If there is absolutely no socially acceptable justification for creating the risk, then

> (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;[110] 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 statistical knowledge 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.[111] Someone who plays involuntary Russian Roulette with a 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.[112] Such an actor

---

[110] But not always. The sadist who wants to badly injure but not kill the victim demonstrates 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).

[111] The Model Penal Code explicitly requires both recklessness as to the risk of death and "extreme indifference." *See* MODEL PENAL CODE § 210.2(1)(b) (1985). Other jurisdictions are less clear about whether recklessness as to the risk must always be proven in order to establish depraved heart murder. *See* WAYNE R. LAFAVE, CRIMINAL LAW §14.4(b), at 783 (5th ed. 2010).

[112] Even if the defendant credibly commits to pulling the trigger only once, many

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.[113]

Accordingly, the *Restatement (Third)*'s criterion for distinguishing individualized from statistical knowledge is essentially sound. In the ordinary case, individualized knowledge should indeed 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."[114] But the criterion must be expanded slightly to account for unusual cases, such as Programmed Rifle or Atmospheric Poison, with a proviso such as the following: "[I]n rare cases, individualized knowledge should extend to 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 unjustifiable."[115] Such a proviso properly treats

---

jurisdictions would uphold a conviction of extreme indifference murder. *See* People v. Roe, 542 N.E.2d 610, 614 (N.Y. 1989); 2 WAYNE R. LAFAVE & AUSTIN W. SCOTT, JR., SUBSTANTIVE CRIMINAL LAW § 7.4, at 202 (1986).

[113] 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 reveals 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.

[114] RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL OR EMOTIONAL HARM § 1 cmt. e (2010).

[115] 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 thus be sanctioned especially harshly – if the actor's reason for creating the risk is sufficiently immoral or unjustifiable. Suppose, although you drive carefully, you take the drive, not for pleasure, convenience, or to accomplish some socially acceptable goal but merely because you sadistically hope that your careful driving will injure or kill someone. *See* ALEXANDER & FERZAN, *supra* note 45, at 50 (discussing a similar example, Frankie). If your improbable plan is realized, my analysis seems to imply that 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 probability of success is very small. But it is much less clear that the careful but sadistic driver should be treated harshly – or indeed should be 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 rarely causes injury unless unusual events or mechanisms intervene. Moreover, if the risks he poses truly are extraordinarily minimal, it is possible that those exposed to them

the culpability displayed in these unusual cases as comparable to the culpability displayed in paradigm individualized knowledge cases.

III. APPLICATIONS

The analysis above is illuminating in a wide range of legal contexts. This section offers a more extended discussion of three such contexts.

A. *Knowledge in Contexts Other Than Knowingly Causing Serious Harm*

"Knowledge" that your action is very likely to cause serious harm is not always culpable. It is only presumptively culpable when you have individualized knowledge. Subject to rare exceptions, only in the individualized knowledge 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?

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) To be liable for murder, a criminal defendant must know that his actions will cause the victim's death.

(2) To be liable for battery, a tort defendant must know that his action will cause offense or physical harm or must know that his action will cause a nonconsensual physical contact.[116]

---

implicitly consent to them or that there is some very modest social benefit to his creating the risks. *See* Simons, *supra* note 99, at 565 (analyzing Alexander's and Ferzan's similar character, Frankie). 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.

On the other hand, if that proof were strong enough, perhaps harsh sanctions for the careful driver who kills are indeed warranted. 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 engaged in a type of behavior that poses only trivial risks *simply* because he wanted to endanger or kill another, perhaps he too should be treated harshly.

[116] Actually, although the requisite "intent" for battery encompasses knowledge, courts differ concerning the requisite *object* of that intent. Some courts require merely an intent to contact another person; other courts also require an intent to harm or offend. Simons, *supra* note 27, at 1066-70.

(3) A business owner must know that he is hiring an undocumented immigrant[117] or a person under the legal age or that he is selling alcohol to a person under the legal age.[118]

(4) A business owner or individual must know that the property that he received from another is stolen.[119]

(5) A lawyer must know that his client will give false testimony.

In each context, the legal 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.[120] Laws prohibiting receipt of stolen property focus on an actor's awareness that a particular item of property is probably stolen, not on his mere awareness that a few of a large number of items in his possession are probably stolen (but he cannot say which ones). In short, some version of the invariant culpability principles most likely applies in all of these cases. Ordinarily, an actor's legal culpability should not change because his business is high volume or his activity extensive.

The five contexts may well differ, however, in just where we should draw the individualized/statistical knowledge line. In this respect, the first two contexts differ significantly from the other three. In the first two contexts, the wrong is serious and of a fundamental sort. A moral wrong underlies the legal wrong,[121] and SSP helps explain why individual knowledge in such cases is especially difficult to justify.[122]

---

[117] 8 U.S.C. § 1324a(a)(1)(A) (2006) ("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.").

[118] *See, e.g.*, MASS. GEN. LAWS ch. 138, § 34 (2000) (prohibiting the provision of alcohol to persons under twenty-one years of age if the owner of the premises knowingly or intentionally provides alcohol to such a person).

[119] 18 U.S.C. § 2315 (2006) (stating that anyone who receives goods worth $500 or more that are known to be stolen shall be fined or imprisoned not more than ten years).

[120] *See* MODEL RULES OF PROF'L CONDUCT R. 3.3 cmt. [6] (2010) (describing the general code of conduct that a lawyer should follow if he knows that the client intends to testify falsely). Research has not disclosed any judicial definitions of "knowledge" in this context that refer to probabilities.

[121] In criminal law terminology, it is *malum in se* (wrong in itself), not *malum prohibitum* (wrong because prohibited).

[122] The second example, however, the tort of battery, encompasses relatively minor harms as well as more serious injuries, because it imposes liability for 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 contexts three, four, and five.

Moreover, one of the reasons for clearly distinguishing statistical from individualized knowledge relates to the grading function of the latter relative to other available grades of an offense, such as recklessness or negligence. But many legal contexts employing a

In the other three contexts, a serious moral wrong need not underlie the legal wrong. Moreover, the specified legal wrong might be a proxy, in useful rule-like form, for the legal or moral wrong that the legal standard is meant to address. Consider the third context, a law against selling alcohol to a minor. This criminal law rule is a pragmatic proxy for an underlying norm: "Don't endanger the health of immature people who do not fully appreciate that they are jeopardizing their own health and safety."

These distinctions between the first two categories and the other categories matter for several 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 not to impose 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 an owner's property will come into the possession of a person with no prior right to it. 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.

Consider the following example:

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. These distinctions are not, however, readily justified by SSP, i.e., by the concern that an actor not impose highly concentrated risks on individual victims. In all three cases, of course, 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 an owner of her property yet does not apply when one creates a 30% or 10% or 1% risk.

Second, in cases such as the third context where a rule serves as a proxy, determining whether individualized knowledge or some lesser culpability should suffice is more complex than in the "knowingly causing serious harm" cases canvassed earlier in this Article, because sometimes a lesser culpability is perfectly consistent with the underlying norm.[123] Third, in a particular legal

---

knowledge requirement do not recognize a lesser grade or degree of culpability. (This is frequently true of contexts three, four, and five.)

[123] For example, requiring only reckless awareness of a risk that the person purchasing

context, we might have good reason to employ a tighter definition or instead a looser one. For example, if it is important to ensure that the insurance exclusion from worker's compensation for "intentional" torts remains either narrow or broad, then we have ample reason to define the individualized knowledge prong of "intentional" torts either narrowly or broadly.[124]

Nonetheless, for practical reasons, legal standards in different domains (e.g., criminal law, tort, regulation) and in different subdomains (e.g., criminal homicide, 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 understandably prefer to use as uniform a criterion as is practicable in different contexts.

B. *If a Cost-Benefit Analysis Supplies Statistical Knowledge, Is the Actor Culpable for Proceeding?*

The propriety and legitimate scope of cost-benefit analysis, whether conducted by government regulators or by private actors, is an enormous topic.[125] Here, I focus on a narrower question. 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 has six parts.

First, to presume culpability in this context is unwarranted. The invariant culpability principles (ICAA and ICREA) tell us that expanding the scope or duration of an activity does not, for that reason alone, increase one's culpability. If we did treat an actor's prediction of harm derived from a thorough cost-benefit analysis as individualized knowledge, the result would be a dramatic and unacceptable expansion of legal liability, 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.[126]

---

alcohol is under age might be consistent with properly respecting the underlying norm not to contribute to the health problems of an immature person.

[124] Some jurisdictions, indeed, define this exception so narrowly that only purpose, not individualized knowledge in any form, satisfies the exception. *See* RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL OR EMOTIONAL HARM § 1 cmt. e (2010). 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. Metro. Life Ins. Co., 970 F.2d 360, 362-63 (7th Cir. 1992).

[125] Useful discussions include MATTHEW D. ADLER & ERIC A. POSNER, NEW FOUNDATIONS OF COST-BENEFIT ANALYSIS (2006); COST-BENEFIT ANALYSIS: LEGAL, ECONOMIC, AND PHILOSOPHICAL PERSPECTIVES (Matthew D. Adler & Eric A. Posner eds., 2001); RISK: PHILOSOPHICAL PERSPECTIVES (Tim Lewens ed., 2007).

[126] If marketing a product in the face of a cost-benefit analysis predicting serious injuries

Second, in this context, heuristic and cognitive biases help explain why people mischaracterize statistical knowledge cases as individualized knowledge cases.[127] Thus, when the causal effect of which the observer is cognizant is a serious harm, people often incorrectly characterize the harm as purposeful, rather than knowing;[128] and, of course, purposeful harm is almost always unjustifiable. Furthermore, psychological research demonstrates that when harms are "on screen" and highly salient, people often conclude that one has an especially strong duty to avoid those harms. This psychological fact, however, is too thin a reed to support a more stringent moral and legal duty.

Third, some versions of cost-benefit analysis are indeed morally and legally objectionable. But even when they are, 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, by itself, is what *makes* acting on the basis of that analysis objectionable. If such an analysis is objectionable, it is for different reasons – for example, 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.

Recall the characters 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. As we have seen, deontologists need not and should not condemn 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 the risks are fairly distributed.[129]

---

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." *See supra* note 28 and accompanying text. Securing consumer benefits such as lower product cost or greater ease of use is not likely to count as a "necessity" justifying a knowing killing. *See* Heinzerling, *Knowing Killing*, *supra* note 11, at 529 ("[E]ven where the necessity defense is recognized in cases of homicide, it is allowed only when an occasion suddenly arises in which the deaths of a few are required to prevent the deaths of many and not when killing becomes a part of long-term social or economic planning.").

[127] *See infra* APPENDIX.

[128] *See* Joshua Knobe, *Intentional Action and Side Effects in Ordinary Language*, 63 ANALYSIS 190, 192-93 (2003).

[129] *See* Simons, *supra* note 16, at 1194 n.54 (distinguishing how deontologists and utilitarians will conduct marginal analyses).

Fourth, it is easy to confuse the highly culpable, cold-blooded, deliberate state of mind that qualifies as first-degree murder with the deliberate, calculating state of mind of a corporate manager or government planner who moves 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."[130] But this conflation of premeditated murder with "premeditated" use of cost-benefit analysis is quite unwarranted; the two categories are not directly comparable. To be sure, if D1 impulsively shoots a victim 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. In this context, impulsiveness sometimes mitigates, and premeditation sometimes aggravates. 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 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 or culpable manner, the question remains: how culpable? The answer depends on what was defective about her analysis. A misunderstanding of the facts? A failure to conduct sufficient research? A failure to consider or 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 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 to fail, because of indifference, even to investigate the likely consequences of its decision to market a potentially dangerous product than to conduct a cost-benefit study of those consequences but perform the analysis negligently.[131] In doctrinal terms, sometimes "recklessness" – in the sense of

---

[130] *See* Heinzerling, *Knowing Killing*, *supra* note 11, at 523, 526. On the other hand, if a product manufacturer can confidently predict that its product will contain some manufacturing defects that will cause harm, perhaps such a prediction helps justify *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 justifiably should impose.

[131] *But see* Viscusi, *supra* note 13, at 552-59 (providing empirical evidence that many people do not appreciate this point; study participants tended to treat a negligently performed cost-benefit analysis as warranting much greater punitive damages than a failure to perform an analysis at all).

awareness that the product will probably cause users a particular aggregate amount of harm – is less culpable than "negligence" – in the sense of blameworthy inadvertence to whether the product will produce such harm.[132] This 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.[133]   But for our purposes, the lesson 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 causes a harm with individualized knowledge.

Sixth, and finally, consider the phenomenon of jury outrage in some high-profile product liability cases – outrage that prompted the award of punitive damages and that appears to have been triggered by the corporation's decision to trade off health and safety against the cost of a precaution.[134] Such cases seem 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. The jury imposed an enormous punitive damage award, and Ford was even criminally prosecuted. A popular understanding of the reason for the award and 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

---

[132] *See* 1 GEORGE FLETCHER, THE GRAMMAR OF CRIMINAL LAW: AMERICAN, COMPARATIVE, AND INTERNATIONAL 315 (2007).

> 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 314-15.

[133] In my view, acting with a cognitive mental state is sometimes more, but 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* Simons, *supra* note 22.

[134] For a recitation of some such cases, see Viscusi, *supra* note 13, at 596-97.

profits it would suffer if it had to incur the cost (which was very low per automobile) of taking a precaution to prevent those harms.[135]

This popular understanding rests on a contestable reading of the facts. It is not clear that the jury that imposed the punitive award and the judges who upheld it decided as they did *simply* because Ford performed a cost-benefit analysis of whether to make the gas tank safer and thus (in their view) improperly traded lives off against money.[136] But I concede that an abhorrence of cost-benefit balancing might indeed have been their central motivation.[137] If that was so, and Ford was sanctioned with punitive damages merely for conducting a cost-benefit analysis, without regard to how sensitively and carefully it performed that analysis, then the punitive damage award was unjustified.[138]

---

[135] *See* Schwartz, *supra* note 13, at 1020-21.

[136] *See id.* at 1039 n.106 (noting that the civil jury in *Grimshaw* was instructed under the consumer expectation test, not a risk-benefit test); *id.* at 1059 n.178 (noting that the enormous size of the 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. That figure, however, was supplied to Ford by a government agency. *See id.* at 1022-26. Also, the famous "smoking gun" cost-benefit memo in which Ford purportedly memorialized its cost-benefit analysis was not admitted into evidence. *Id.* at 1021.

[137] *See id.* at 1067. In *Grimshaw v. Ford Motor Co.*, 174 Cal. Rptr. 348 (Cal. Ct. 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.

Id. at 384. The court went on to state,

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

Id. at 385 (emphasis added). These excerpts, however, can also be read as expressing a more radical view, that any use of cost-benefit analysis in designing products is unjustifiable and warrants punitive damages.

[138] 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. Yet murder (which under Indiana law encompasses

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 justify 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:

   • 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 retributivists' view, is that consequentialists are committed to supporting the intentional punishment of identified innocents if this would serve a greater good. For example, consequentialists would support the lynching of an innocent if this were necessary in order to appease a mob and prevent much greater violence.

   • Consequentialists reply that the retributivists are not on such high moral ground, because every practicable, real-world system of criminal justice will inevitably make at least a few mistakes in attempting to identify and punish 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, they seem to be guilty of a fundamental inconsistency.   On the one hand, they condemn

---

knowingly causing a death) would have been an appropriate charge if indeed the law treats statistical knowledge the same as individualized knowledge.   Indiana homicide law applicable to the 1978 crash defined murder as including knowingly causing death. *See* IND. CODE § 35-42-1-1 (2004).

   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 allegedly reckless initial defective design. Schwartz, *supra* note 13, at 1017.

consequentialists for sometimes permitting the greater good to justify the intentional punishment of the innocent. 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 invoke the overriding importance of a greater good – protecting the rights of victims or giving offenders their just deserts – to justify a lesser evil – the use of a punishment system that we know will inevitably punish a small number of innocents.[139]

The earlier analysis 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 utilitarian judge decides to punish an individual defendant he knows is factually innocent. This is both illegal and immoral, even if the judge reasonably believes that such punishment will provide substantial social benefits.[140] Such an action is indeed analogous to the lynching example just mentioned. But when a judge decides to participate in a system that incorporates a reasonable doubt standard and that he knows will therefore convict a very small number of innocent persons,[141]

---

[139] *See* Russell Christopher, *Deterring Retributivism: The Injustice of "Just Punishment*," 96 NW. U. L. REV. 843, 899-90 (2002) ("Retributivism justifies the knowing or intentional implementation and utilization of necessarily fallible legal institutions of punishment in which it is foreseeable that some unknown innocents will inevitably be mistakenly punished. By justifying such fallible institutions, retributivism is knowingly punishing the innocent."); George Schedler, *Can Retributivists Support Legal Punishment?*, 63 MONIST 185, 185 (1980).

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

Some try to rebut the "knowing punishment of the innocent" 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. *See* MICHAEL MOORE, PLACING BLAME: A GENERAL THEORY OF CRIMINAL LAW 158 (1997).

But this move fails. The distinction between purpose and knowledge is not an adequate explanation. In an individual case, if the judge knowingly but not purposely punishes an innocent, his action is still almost impossible to justify. Thus, in the utilitarian judge example, the judge might have acted with individualized knowledge of the defendant's innocence, but not with purpose to punish an innocent. If his overriding purpose was to protect the public and 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.

[141] 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.

his action is neither illegal nor immoral. The difference is based on the distinction between individualized and statistical knowledge.[142]

In virtually[143] no case is it morally justifiable to impose criminal punishment on an individual criminal defendant whom the legal authorities know is factually innocent. For the reasons we have canvassed, however, it does not follow 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 ICAA it is also permissible to apply that standard in an indefinitely large number of cases. (None of the exceptions to or qualifications of ICAA applies 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 an innocent person. For example, she employs a jury instruction that is insufficiently stringent and knows that the instruction 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 both morally and legally, it is less wrongful than upholding the conviction of an individual defendant who she knows is very likely innocent.

This analysis is subject to two qualifications. First, even when a legal actor neither harbors a purpose to punish an innocent person as such[144] nor possesses

---

[142] The relevant distinction is not between "known" and "unknown" victims. *See* David Dolinko, *Three Mistakes of Retributivism*, 39 UCLA L. REV. 1623, 1633 (1992); Richard O. Lempert, *Desert and Deterrence: An Assessment of the Moral Bases of the Case for Capital Punishment*, 79 MICH. L. REV. 1177, 1183, 1226 (1981). Rather, the distinction 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 also infra* APPENDIX.B (discussing the identifiability of the victim).

[143] 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 September 11, 2001.

[144] 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, 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

individualized knowledge that the person being punished is innocent, her statistical knowledge that innocent individuals are being punished is still morally and legally very significant. A system that leads to such results demands a very compelling justification, and actors within the system have a stringent duty to minimize the risk of innocent individuals being punished, especially when they are aware of the scope of that risk.

Indeed, a punishment system itself can be illegitimate when it permits too many wrongful convictions. 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 of convicting the innocent is unjustifiable, given the irrevocability and seriousness of the death penalty.

In a provocative article, Cass Sunstein and Adrian Vermeule explicitly claim that the risk of mistake does not distinguish the death penalty from other social policies that predictably and inevitably cause or permit the deaths of innocents.[145] Many would reject their claim, believing that the statistically certain death of even one innocent individual due to the state's death penalty system is so grievous a wrong that no consequential benefits can justify it.[146] 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 death, their insistence that human institutions inevitably produce some level of error[147] is a valid point worth remembering. Furthermore, an infinitesimal risk of error should not necessarily condemn a capital punishment system if the system 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.[148]

Second, the State should, of course, provide an ex post remedy if a case of mistaken punishment of the innocent is discovered. But whether the system

---

person as such. And the latter type of intent will often be lacking. Thus, an actor lacks a genuine intent to punish an innocent if, although he intends to *punish* X, knowing that X is innocent, nevertheless the *innocence* of X is not even part of the reason that he intends to punish him. Suppose a 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 would strongly prefer that X were not innocent and indeed hopes that X might actually be guilty. The judge does not intend to punish an innocent, but his knowledge that he is punishing an innocent is still highly culpable.

[145] *See* Cass R. Sunstein & Adrian Vermeule, *Is Capital Punishment Morally Required? Acts, Omissions, and Life-Life Tradeoffs*, 58 STAN. L. REV. 703, 721 (2005) (arguing that the death penalty is morally obligatory if it deters a significant number of private killings).

[146] *See, e.g.*, Michael L. Radelet & Hugo Adam Bedau, *Fallibility and Finality: Type II Errors and Capital Punishment*, *in* CHALLENGING CAPITAL PUNISHMENT 91, 106 (Kenneth C. Haas & James A. Inciardi eds., 1988).

[147] Sunstein & Vermeule, *supra* note 145, at 728.

[148] *See* Steiker, *supra* note 78, at 784.

that results in such mistakes is unjustifiable and whether individual actors within the system bear moral or legal responsibility for such mistakes are separate questions. The answer to these questions, I suggest, is "no" – or more precisely, "not necessarily." We should be careful not to answer "yes" to these questions based only on our unthinking conflation of statistical with individualized knowledge.

D.   *Practical and Doctrinal Issues*

I do not focus in this Article on practical and doctrinal issues. No doubt, significant difficulties of proof arise if legal sanctions depend on the actor's specific beliefs about the precise probability that she will cause harm or that a legally relevant circumstance exists. In response, the law has employed various strategies. 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 that expand the meaning of "knowledge" to encompass lower-probability beliefs.[149] 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, from the objective physical acts themselves, an actor's purpose in so acting than to determine what he knows about the circumstances of the act or about its possible results.[150]

Moreover, pragmatic concerns also shape the categories that the law employs: it must draw a somewhat arbitrary line between individualized knowledge and cognitive recklessness, distinguishing individualized knowledge by a rough criterion of probability, such as "practical certainty" or "highly probable" or "more likely than not."[151]

Although the analysis in the Article 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

---

[149] Consider the doctrine of willful blindness, which essentially treats certain unusually culpable instances of reckless belief as satisfying a legal knowledge requirement. *See* Simons, *supra* note 99, at 561.

[150] After all, those types of knowledge will often have no direct motivational significance and thus might not produce a distinctive pattern of conduct that provides strong evidence of his state of mind.

[151] For a recent analysis suggesting that juries might have great difficulty distinguishing knowledge from recklessness, see Francis X. Shen et al., *Sorting Guilty Minds*, 86 N.Y.U. L. REV. 1306, 1337-1344 (2011). 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., a "substantial," "significant," or "nontrivial" risk.

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 from whether a repetition of similar acts justifies a recidivist premium in criminal law.

<div align="center">CONCLUSION</div>

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 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.

- ***Statistical victims are still victims.*** 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.[152] A primary explanation for this position is SSP, the deontological principle that treats an actor as highly culpable when he knowingly imposes a highly concentrated risk of serious harm on a victim.

- ***Juries often reject the distinction between statistical and individualized knowledge.*** This behavior is a relevant datum. 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. Often, however, when juries reject the distinction, they are wrong to do so; they are often improperly neglecting the invariant culpability principles or are misled by cognitive biases or psychological heuristics.

---

[152] By contrast, when we are considering ex ante what 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 twenty murders. *See supra* text accompanying note 77.

- **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.

- **Acting with merely statistical knowledge is sometimes still highly culpable.** This is correct, but only in limited circumstances. The most prominent such circumstance is when the risk is extremely unjustifiable.[153]

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

- **Scale should not matter.** This proposition 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 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.

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

- **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.   In this context, however, it usually provides the right answer.

- **Knowingly harming another is especially wrongful only when the actor's knowledge is individualized.** This argument is indeed crucial, but we must distinguish two versions of it:

  - *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.

---

[153] Recall the hypotheticals Time Bomb and Atmospheric Poison, but contrast Atmospheric Cancer Cure.

o *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. We should distinguish, however, 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 person to whom the actor sold alcohol is underage. It is doubtful that SSP applies in the latter context.

The analysis in this Article also has other important implications:

- 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. This special type of culpability justifies greater legal sanctions than the invariant culpability principles would permit. 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."[154] 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, his 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, to a statistical certainty, punish some innocent defendants. 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

---

[154] RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL OR EMOTIONAL HARM § 1 cmt. e (2010).

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 questions about the conundrum of statistical knowledge.[155] Nevertheless, the Article does provide a novel analytic structure, one that should assist future voyagers who wish to investigate this important, yet underexplored, region of the culpability terrain.

---

[155] 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. *See supra* text accompanying notes 75-81.

APPENDIX: HEURISTIC ERRORS AND VICTIM IDENTIFIABILITY

The intuitive reactions of ordinary people to particular illustrations of statistical and individualized knowledge are worthy of respect. At the same time, psychological research reveals the existence of numerous heuristic errors and cognitive biases that could affect such reactions and thus could undermine the weight that they deserve. Here, I briefly review some of the relevant concepts and literature.

A.   *The Significance of Large Numbers*

The tendency to view statistical knowledge as very similar, if not identical, to individualized knowledge might reflect a specific heuristic error.[156] Many people have difficulty apprehending the significance of large numbers, and they therefore mistakenly assume that if an activity will predictably 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.[157]

Sunstein concludes that the response he predicts to these examples illustrates a "moral heuristic," a rule of thumb (analogous to a cognitive heuristic) that works well most of the time but also sometimes misfires[158]:

---

[156] *See* Cass Sunstein, *Moral Heuristics and Moral Framing*, 88 MINN. L. REV. 1556, 1578-80 (2004); Viscusi, *supra* note 13, at 574 (describing a products liability case in which a key witness presented evidence of the automobile manufacturer's conscious cost-benefit analysis; the jury awarded punitive damages, and "they knew" became the "'constant refrain among the jurors interviewed'" (quoting Terence Moran, *GM Burns Itself*, AM. LAW., Apr. 1993, at 68, 83)).

[157] Sunstein*, supra* note 156, 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 example were phrased less dramatically: "Company A knows that its product will cause the death of ten people," rather than "will kill ten people."

[158] *See id.* at 1578.

> 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.[159]

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. A plaintiff's trial lawyer would of course prefer the rhetoric of the Company A example rather than that of the Company B example, because she will want to suggest to jurors that their (justifiable) disapproval of actions taken with *individual* knowledge should cause them to disapprove of actions taken with only statistical knowledge.[160] Nevertheless, 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 risks, at least relative to more salient and more immediate risks and harms.[161]

B. *Is the Real Question Whether the Victim Is Identifiable?*

The 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

---

[159] *Id.* at 1579-80.

[160] 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 the number of accidents or injuries alone. *See, e.g.*, Loitz v. Remington Arms Co., 563 N.E.2d 397, 404-07 (Ill. 1990) ("The 94 prior accidents represent about 0.003% of Remington's production of more than 3 million Model 1100 barrels during the span of time in question. The prior accidents must be considered not only in relation to the number of barrels produced . . . but also in relation to the estimated number of times those guns would have been used."); Kopczick v. Hobart Corp., 721 N.E.2d 769, 776 (Ill. App. Ct. 1999) (pointing out that the thirty 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.").

[161] *See* Albert W. Dzur & Alan Wertheimer, *Forgiveness and Public Deliberation: The Practice of Restorative Justice*, 21 CRIM. JUST. ETHICS 3, 14-15 (2002) (illustrating the human tendency to weigh the interests of identifiable lives much more heavily than the interests of future statistical lives). Often people discount very heavily, or simply ignore, low probability risks, a cognitive bias that is classified as "probability neglect." *See* Cass R. Sunstein, *Probability Neglect: Emotions, Worst Cases, and Law*, 112 YALE L.J. 61, 62 (2002).

is oversimplified and, 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 that (1) the strong intuition that "identifiability" is morally and legally highly significant can largely be explained by psychological heuristics and cognitive biases that lack direct moral relevance, and (2) the intuition, nonetheless, 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, but when the actor 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 society[162] 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 a currently unidentifiable miner in the future.[163] Moreover, legal doctrine contains some examples of especially stringent duties to rescue identifiable individuals. First, under traditional tort doctrine, a landowner owes a duty to adult trespassers only if he "knows" of their presence; his duty is much more limited if he is aware only of a statistical risk of trespassing.[164] Second, courts

---

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

[163] *See* CALABRESI, *supra* note 11, at 6; Fried, *supra* note 78, at 1415, 1416. The classic quotation on the topic is from Josef Stalin: "The death of a single Russian soldier is a tragedy. The death of a million soldiers is a statistic." *See* RICHARD E. NISBETT & LEE ROSS, HUMAN INFERENCE: STRATEGIES AND SHORTCOMINGS OF SOCIAL JUDGMENT 43 (1980).

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's or hospital's care for an existing ailment, a more stringent duty is owed to the patient to prevent further harm. Specifically, greater resources are required to be expended to prevent harm from the existing ailment 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, *supra* note 12, at 46; *see also* Phaedra S. Corso et al., *Assessing Preferences for Prevention Versus Treatment Using Willingness to Pay*, 22 MED. DECIS. MAKING, S92, S92 (2002) (finding that survey respondents stated a willingness to pay for treatment that was significantly greater than their willingness to pay for prevention); Fried, *supra* note 78, at 1416 ("A developing nation may devote large resources to general medical facilities, when public health measures would reduce the death rate more efficiently."); Karen E. Jenni & George Loewenstein, *Explaining the "Identifiable Victim Effect*,*" 14 J. RISK & UNCERTAINTY 235, 236 (1997) (explaining how the bias toward identifiable victims has manifested itself in the national healthcare debate in that "expensive measures are often taken for identified individuals, but funding for preventative care seems to be lacking").

[164] *Compare* RESTATEMENT (SECOND) OF TORTS § 336 (1965) ("A possessor of land who knows or has reason to know of the presence of another who is trespassing on the land is

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 the plaintiff cannot show that but for the breach of duty, the decedent more likely than not would have been saved.[165]

The significance of identifiability, however, is a more complex issue than first appears, because identifiability has various distinct senses.[166] 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; it is therefore difficult to isolate 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; in discussing psychological research however, I will necessarily have to take the view of ordinary observers, since that is the perspective that the studies explore.[167]

---

subject to liability for physical harm thereafter caused to the trespasser by the possessor's failure to carry on his activities upon the land with reasonable care for the trespasser's safety."), *with id.* § 333 cmt. a (explaining that a landowner is not liable to a trespasser for physical harm caused by his failure to exercise reasonable care "although from past experience or otherwise the possessor has every reason to realize that there is a strong probability that trespassers will intrude upon his land").

[165] *See* Gardner v. Nat'l Bulk Carriers, Inc., 310 F.2d 284, 287 (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." RESTATEMENT (THIRD) OF TORTS: LIAB. FOR PHYSICAL AND EMOTIONAL HARM § 41 (2010). But this limitation is due to the perceived unfairness of imposing broad liability on therapists and the perceived inefficacy of such a broad duty in reducing harm. It is not premised on the principles we have discussed, e.g., the supposed greater culpability of one who harms a named or non-anonymous victim, or the special duty to avoid imposing a concentrated risk on an individual.

[166] For a careful parsing of the different senses and sources of the identifiable victim effect in the psychological literature, see Jenni & Loewenstein, *supra* note 163, at 237-38 (discussing four differences between identifiable victims and statistical victims which may account for the identifiable victim effect: (1) vividness, (2) certainty and uncertainty, (3) proportion of the reference group that can be saved, and (4) ex post versus ex ante evaluation). *See also* CASS SUNSTEIN, *Cognition and Cost-Benefit Analysis*, *in* COST-BENEFIT ANALYSIS: LEGAL, ECONOMIC, AND PHILOSOPHICAL PERSPECTIVES 223 (2001) (providing a helpful overview of cognitive biases that affect risk perception).

[167] Another relevant perspective is that of an idealized observer. 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 versus ex ante in

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, age, or name, or might be personally acquainted with the victim.

Second, identifiability could mean *salience*: the risk to the victim is vivid and emotionally engaging.[168] 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 is at risk of suffering health and safety consequences due to inadequate advanced health and safety precautions in the mine, these miners are virtually invisible to public attention.[169] The availability heuristic is likely playing a role here: people view a type of event as more probable if it more easily comes to mind.[170] And the emotional connection to a victim is strongest when only a single victim is "on screen."[171]

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. 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.[172] Psychological studies confirm that people react more generously

---

prevention.

[168] Jenni & Loewenstein, *supra* note 163, at 237.

[169] 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; but it is almost always clear, in an individualized knowledge case, that the actor caused the death.

[170] *See, e.g.*, Sunstein, *supra* note 156, at 229-30. 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, *supra* note 163, at 253.

[171] 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." Tehila Kogut & Ilana Ritov, *The "Identified Victim" Effect: An Identified Group, or Just a Single Individual?*, 18 J. BEHAV. DECISION MAKING 157, 165 (2005).

[172] 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 suffer from a mine collapse if we do not act now to build more safety

when they confront someone already in peril than someone who is expected to be in peril in the future.[173]

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 do not 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 a trapped miner will die if not rescued) 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).[174]

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 saved in the future by a current precaution are not. Just as we feel stronger bonds to those in our geographic community, we often feel a stronger obligation to a person whose plight is immediately before us.[175]

---

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" while the second miner is not. *See* Fried, *supra* note 78, at 1419. But this idea of "actual peril" is difficult to cash out, at least if one views risks as subjective and epistemic.

[173] *See* Deborah A. Small & George Loewenstein, *Helping a Victim or Helping the Victim: Altruism and Identifiability*, 26 J. RISK & UNCERTAINTY 5, 6 (2003) (observing a difference in attitudes towards ex post and ex ante precautions). Strikingly, psychological research reveals that people 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. *See id.* at 10. 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. *See id.* at 5-14.

[174] According to Jenni & Loewenstein,

[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, *supra* note 163, at 238.

[175] A recent study finds: "Knowing that someone will be hurt, but not knowing *when* that person will be hurt, does reduce judgments of intentionality." Pam Mueller, Lawrence M. Solan & John M. Darley, When Does Knowledge Become Intent?: Perceiving the Mind of Wrongdoers 3 (October 22, 2011) (unpublished manuscript), *available at* http://papers. ssrn.com/sol3/papers.cfm?abstract_id=1884518. Surprisingly, this particular study finds that some other aspects of identifiability make no difference to culpability judgments. For

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.[176] 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 the seven different senses of identifiability in turn.

First, consider non-anonymity. If identifiability is understood in its very weak sense, as merely requiring some information that makes the victim non-anonymous, it 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 – for example, "future coal miners who might benefit from a mine safety improvement." Even in the stronger sense, requiring significantly more information, the concept is an unhelpful criterion of individualized knowledge. A driver speeding at nighttime might perceive only the dark shape of a human body that he knows his car is about to hit. 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 know that one person is in need of rescue, but we are unable to identify him 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

---

example, it finds that "[k]nowing that someone will be hurt, but not knowing *who* will be hurt, has no effect on [culpability judgments]." *Id.* But the authors plausibly surmise that this null effect might be due to the small size of the population at risk in the scenarios. *See id.* at 26-27 ("[P]articipants may have seen [the population at risk] as a small enough group to constitute an 'identifiable victim' en masse.").

Similarly, a familiar criticism of consequentialism is the "nearest and dearest" objection, i.e., the complaint that consequentialists cannot account for a perceived greater duty to our neighbors or fellow citizens. *See* KAMM, INTRICATE ETHICS, *supra* note 74, at 345-97 (defending a complex view concerning how proximity bears on our deontological moral obligations); Frank Jackson, *Decision-Theoretic Consequentialism and the Nearest and Dearest Objection*, 101 ETHICS 461, 461 (1991) (suggesting that consequentialism, by ignoring the special duties owed to those nearest and dearest to us, is in conflict with what makes life worth living).

[176] 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 163, at 252.

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 culpable.[177]

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 observers should, without more, have normative weight.[178] Moreover, salience and vividness are not always present in cases of individualized knowledge. Consider again the speeding nighttime driver from the prior paragraph, who knows very little about the potential victim. Nor are salience and vividness always absent in cases of statistical knowledge. Consider 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.[179] 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 immediate peril.

Fourth and fifth, causal and statistical uncertainty do indeed characterize many cases where advanced precautions might have a beneficial effect on safety levels. By contrast, we tend to be more certain 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 is

---

[177] In some cases an actor's willingness to harm an anonymous stranger seems to inculpate, for it makes the act more incomprehensible, more frightening, and more inhuman. *See* Heinzerling, *Statistical People*, *supra* note 11, at 196 (observing that in the 1982 Tylenol poisoning incident, it was the 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 inculpate, for it can express personal spite or malice or can involve targeting the victim's vulnerability in an especially reprehensible way.

[178] A policy that gives higher priority to saving more familiar victims "would amount to allowing media coverage to determine aid allocation." Jenni & Loewenstein, *supra* note 163, at 240; *see also* Kenneth W. Simons, *Dworkin's Two Principles of Dignity: An Unsatisfactory Nonconsequentialist Account of Interpersonal Moral Duties*, 90 B.U. L. REV. 715, 721-24 (2010).

[179] *See supra* note 172.

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. Moreover, an actor with individualized knowledge that he is likely to harm another might still harbor small doubts about the causal effects of his act.

Sixth, temporal and spatial proximity do not fully explain the coal miner examples; many Americans would empathize with victims of a mine disaster in England as strongly as with victims in another locality within the United States. And these features do not align very well with individualized knowledge; modern technology permits actors to knowingly harm others 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. It also might help explain some of the psychological findings noted above.[180] In the "trapped coal miner" scenario, we do have a concentrated risk – that is, a very high probability that this individual will die if not rescued. By definition, then, the person who fails to save him often has *individualized knowledge* that the miner will die. Moreover, SSP 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 have discussed, including Alfa, 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.[181]

---

[180] Indeed, Jenni and Loewenstein find that of the various factors that might explain the identifiable victim effect, one of the strongest was the proportion of a reference group that could be saved. *See* Jenni & Loewenstein, *supra* note 163, at 254-55 (finding that the factors most critical to the strong response elicited by the identifiable victim were the certainty that the victim would die without aid and the fact that the victim comprised 100% of the risk group). 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 evaluating the effect is susceptible to quite arbitrary framing in a way that SSP is not. *See id.* at 254 ("The reference group is often largely a matter of framing, and it is difficult to defend a distinction between a situation where there is a group of 10 randomly distributed 'vaccine sensitive' people who are at risk of death from a flu vaccine, but who cannot be identified beforehand, and a situation in which 10 random people will be killed by the same vaccine.").

[181] Many deontologists 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

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 moral significance? Some, for example, endorse this expressive or symbolic argument: in order to uphold our self-conception as a caring society, we have a stronger duty to aid persons in need of rescue when their plight is salient and unmistakable.[182] I am inclined to agree with skeptics who reject this argument. For the argument might simply reflect an irrational psychological disposition to overvalue salient events, an inability to comprehend the importance of ex ante precautions, or an unjustifiable tendency to discount greatly, or even to ignore, risks of future harm when the probabilities of such harm are low.

However we resolve this debate, it is 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 individual wrongdoing. To be sure, that kind of expressive value might sometimes affect which actions are permissible in the first place, especially when the actor in question is the government; the acts and omissions of the government do 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 higher duty. On this approach, it might be impermissible for a coal mine operator not to expend enormous resources to save a trapped miner, even if it would be permissible not to spend equivalent resources in advance on safety precautions to save a statistical life. But I think we should question the premise of this approach, that the symbolic value of rescuing should increase a private actor's

---

not to harm. Thus, SSP would impose a much more stringent duty to aid 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 when the vulnerable person has only a small chance of dying if not aided.

[182] *See* Gibbard, *supra* note 12, at 101 ("It may . . . be dehumanizing to stand idly by when strenuous, expensive effort has a substantial chance of saving lives."); Hammitt & Treich, *supra* note 12, at 46 ("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."); *see also* RONALD DWORKIN, JUSTICE FOR HEDGEHOGS 178-79 (2010) (arguing that it is "wrong to ignore the natural responses that a respect for life provokes when we are directly confronted with a person facing imminent death"); Fried, *supra* note 78, at 1425 (describing but questioning the "symbolic value" argument).

I concede that this 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.

culpability or should render his otherwise permissible acts impermissible. In my view, we should only raise the level of his duty to the height that SSP demands.   Thus, we should sometimes, but not always, require greater expenditures on saving identifiable lives.

DRAMATIS PERSONAE AND CHARTS