University of Pennsylvania Carey Law School

Penn Law: Legal Scholarship Repository

Faculty Scholarship at Penn Law

1998

Incommensurable Choices and the Problem of Moral Ignorance

Leo Katz

University of Pennsylvania Carey Law School

Follow this and additional works at: https://scholarship.law.upenn.edu/faculty_scholarship

Part of the [Criminal Law Commons](#), [Ethics and Political Philosophy Commons](#), [Law and Philosophy Commons](#), [Legal History Commons](#), and the [Philosophy of Mind Commons](#)

Repository Citation

Katz, Leo, "Incommensurable Choices and the Problem of Moral Ignorance" (1998). *Faculty Scholarship at Penn Law*. 1514.

https://scholarship.law.upenn.edu/faculty_scholarship/1514

This Article is brought to you for free and open access by Penn Law: Legal Scholarship Repository. It has been accepted for inclusion in Faculty Scholarship at Penn Law by an authorized administrator of Penn Law: Legal Scholarship Repository. For more information, please contact PennlawIR@law.upenn.edu.

INCOMMENSURABLE CHOICES AND THE PROBLEM OF MORAL IGNORANCE

LEO KATZ[†]

If I cannot decide between *A* and *B*, that would seem to show that I am indifferent between them. But as Joseph Raz famously demonstrated in *The Morality of Freedom*,¹ it does not show that at all. His demonstration basically consisted of pointing to a feature of such situations that previously had been overlooked completely. If I cannot make up my mind between a Toyota and a Honda, it certainly looks as though the reason I am dithering is that I am nearly indifferent between them. But consider, said Raz, that I would have no problem making up my mind between one such Honda and another such Honda that happened to be selling for a few dollars less. If I were truly indifferent between the more expensive Honda and the Toyota, Raz observed, then I should no longer be indifferent between the cheaper Honda and that self-same Toyota: I should prefer the cheaper Honda. But since I find choosing between the cheaper Honda and the Toyota just as hard as choosing between the more expensive Honda and the Toyota, something else must be going on.

What exactly is going on in such a case is, of course, quite mysterious. Hence this Symposium.² Everyone would, I think, agree that what is going on has something to do with the fact that I find the cheaper Honda easy to compare with the more expensive Honda, but that I find either of the Hondas quite hard to compare with a Toyota—they are just so different. They are, it seems, incommensurable. My being unable to decide between the two cars thus seems not to stem from my being in equipoise between them (otherwise a slight drop in the price of the Honda immediately would break the tie), but from my being unable to compare them properly.

[†] Professor of Law, University of Pennsylvania Law School.

¹ JOSEPH RAZ, *THE MORALITY OF FREEDOM* 321-66 (1986).

² I am using "incommensurability" to refer to what some prefer to call "incomparability." See, e.g., Ruth Chang, *Introduction*, in *INCOMMENSURABILITY, INCOMPARABILITY, AND PRACTICAL REASON* 1 (Ruth Chang ed., 1997). I prefer that terminology because it is better at evoking what I argue to be the root of the incommensurability phenomenon in this Article.

In the first Part of this Article, I will try to show that the roots of this mysterious-seeming incommensurability phenomenon are often very mundane and that the problem of what to do about incommensurable choices is therefore often very easy to solve. To be more precise, I will try to show that things have a way of seeming incommensurable for no other reason than that the chooser happens to be fairly uninformed about them. Once his ignorance is dispelled—as a result of nothing more than a bit of sustained investigation and reflection—the incommensurability generally will disappear. In the remainder of the Article, I will pursue some questions that arise out of this: What is one to do if one has not yet been able to come by the information or the insight that would make the incommensurability disappear? And how bad is it if one gets things wrong?

I. HOW IGNORANCE LEADS TO INCOMMENSURABILITY

The idea that incommensurability is simply a manifestation of ignorance is not original with me. Donald Regan has taken a line somewhat like this in his extended essay on Raz's book.³ Here, I will pursue the point farther than he does, as well as pursue the further interesting questions to which it gives rise. The best way to explain what I have in mind is with an example. Suppose we face the task of comparing two irregularly shaped pieces of paper as to their size. One of these pieces of paper resembles in outline the State of Texas, the other the State of Idaho. To look at them, it seems possible that they are equal in area, or that one of them is larger. There seems no easy way to tell which is the case. Suppose now there exists a third piece of paper, which in fact has the same shape as "Idaho," but on a slightly smaller scale. We are now finding ourselves in a situation that has all the earmarks of incommensurability. We are unable to choose as between Texas and Idaho which is larger. We have no trouble deciding that "big" Idaho is larger than "small" Idaho. Nevertheless, this does not mean that we are able to choose between Texas and "small" Idaho as to which is larger. The reason for this perceptual kind of incommensurability seems to be the relative ease with which we can compare the two Idaho shapes and the difficulty of comparing either of those shapes with the Texas shape. Note, however, that the

³ Donald H. Regan, *Authority and Value: Reflections on Raz's Morality of Freedom*, 62 S. CAL. L. REV. 995, 995 (1989) (detailing general agreement with Raz because "[h]e comes closer to the truth about political morality than anyone has for nearly a century," but also elaborating on certain "tensions" in Raz's thought).

incommensurability we have encountered here can be dispelled relatively easily. All we need to do is to stop trying to make the judgment by the eyeballing method and resort to more advanced measuring instruments instead. The incommensurability here is simply the result of ignorance: our inability to do with the naked eye what we could easily do with some measuring tools. Dispel the ignorance and you have eliminated the incommensurability.

In an implicit way, the law has, in fact, long recognized the intimate connection between incommensurability and ignorance. Consider the way the burden-of-proof rules operate in a civil case. Imagine a plaintiff who makes an allegation and backs it up with the meagerest of evidence. The defendant in turn denies it, offering no evidence whatsoever. Under such circumstances, a court is obliged to dismiss the case on the ground that the plaintiff has not met his burden of proof. Let us explore a little bit what it means to say here that the plaintiff has not met his burden of proof. It clearly does not mean that the defendant has by a preponderance of the evidence *disproved* the plaintiff's allegation. Rather, it means that the court has not been provided enough information to be able to decide which of the two sides is right. How should we think about the judge's inability to decide which side is right? Are we to think of him as being in equipoise between the two parties? Does the defendant win because the plaintiff has only managed to put the case into equipoise rather than the "50%-plus-a-smidgen" range that would entitle him to win? The burden-of-proof rules make it clear that the court's inability to choose between the two sides does not show it to be in equipoise: If such were the case, then the plaintiff's introduction of a mere scintilla of additional evidence would entitle him to a favorable judgment. It is quite clear, however, that much more is demanded of him. The burden-of-proof rules thus treat the judge's inability to decide as a sign of incommensurability rather than indifference.

In probability theory too, the connection between incommensurability and ignorance long has been recognized. Somebody asserts a proposition, the truth of which I have no idea. I am unable—to put the matter with pedantic, but I think illuminating, clumsiness—to decide between two alternatives: (1) the proposition is true; and (2) the proposition is false. Does my inability to choose reflect that, given how little information I have on the matter, each alternative is equally likely? Although it is tempting to think so, that would lead to contradictions. The root cause of those contradictions turns out to be the

fact that my inability to choose between the two alternatives reflects not "indifference" but "incommensurability."

Consider the statement "Brian will win his next gamble." Not having any evidence whatsoever on the matter—not even knowing what kind of gamble is being contemplated—you might be inclined to consider yourself "indifferent" on the issue and assign a 50% probability to the possibility that Brian will win, and a 50% probability to the possibility that Brian will lose. Consider next the statement "Brian will win a million dollars in his next gamble." This statement, being more specific than the first statement, is less probable than the first statement. Nevertheless, since you have no evidence on the matter, you should consider yourself "indifferent" on the issue and assign a 50% probability to the possibility that Brian will win a million, and a 50% probability to the possibility that Brian will not. The result, of course, is that you end up in a contradiction. Because, in fact, you are not indifferent between the two possibilities, you are not truly in equipoise. The two alternatives are simply incommensurable.

If I am right about the connection between incommensurability and ignorance, then there should be many cases in which previously incommensurable alternatives are rendered commensurable through the influx of additional information or insight. And indeed, there is no dearth of such cases. Every time a plaintiff backs his complaint up with evidence, he turns an incommensurable choice into a commensurable one. But those are admittedly very mundane, factual kinds of choices. What about incommensurable choices that have a heavier moral tinge to them? Are there examples here too of incommensurability being dispelled through insight or information? Let me suggest a few.

A. Utility Measurement

Economists and utilitarians used to be much bothered by a very basic kind of incommensurability: our inability to compare the happiness derived from our first one thousand dollars with the happiness derived from our last one thousand dollars. Intuitively, it seemed right to say that the latter conferred much less joy than the former. But no one felt that he had a truly secure grip on that comparison.

That insecurity was especially apparent when one was asked to compare, say, the joy derived from having one's income jump from \$10,000 to \$20,000 with the joy derived from having one's income jump from \$100,000 to \$200,000. Those two kinds of joy seemed

truly incommensurable. It was clear that the joy of seeing one's income rise from \$100,000 to \$200,001 was greater than seeing it rise from \$100,000 to \$200,000. Yet it did not therefore follow that that joy was greater than the joy of seeing one's income rise from \$10,000 to \$20,000.

In the end, it turned out that our inability to compare these different quantities of happiness was much like our inability to compare the size of two irregularly shaped pieces of paper. They are incommensurable if we insist on settling the matter by eyeballing. They become commensurable once we resort to more refined technology. In the case of utility measurement, that technology was the Von Neumann-Morgenstern approach to constructing a cardinal utility function. By making the simple but very ingenious assumption that everyone is trying to maximize his (statistically) expected utility (regardless, that is, of its statistical distribution), and then asking people to express their preferences among various easy-to-compare gambles, Von Neumann and Morgenstern were able definitively to answer questions thought to be inherently unanswerable. They were able to compare things that had been incommensurable, and were now actually in a position to say whether, for a given person, the jump from \$10,000 to \$20,000 represented a greater or lesser increase in happiness than the jump from \$100,000 to \$200,000.

B. *The Concept of Desert*

Incommensurability quickly rears its head when we talk about desert. Is someone receiving more or less recognition than he deserves? The incommensurability here seems so great that many think the question vacuous. Here too, however, one can eliminate at least some of the incommensurability (and with sustained effort maybe all of it) by casting about for some appropriate measurement tools that would allow us to do more than just "eyeball" the question.

Those measurement tools can be found in the criminal law, which provides us with a highly developed set of rules for measuring "negative desert," that is, blameworthiness. With just a little bit of effort, those rules can be converted into a means of measuring "positive desert," that is, praiseworthiness. The best way to demonstrate this is to juxtapose the sorts of things we say about a group of thugs committing a murder with what we say about a group of soldiers killing for a just cause. The things we say turn out to be remarkably symmetrical. About the thugs we say that those among them who bring about harm by an act are worse than those who bring it about

by an omission; those who bring it about intentionally are worse than those who bring it about recklessly; those who bring it about proximately are worse than those who bring it about more indirectly; those who bring it about as a principal are worse than those who merely act as an accomplice; those who complete the crime are worse than those who merely attempt it. All of this would seem to apply equally when we talk about soldiers fighting for a just cause. Those among them who bring about a good consequence by an act are more praiseworthy than those who merely do it by an omission, that is, by not stopping others from doing so; those who bring it about intentionally are better than those who bring it about inadvertently; those who bring it about proximately are better than those who do so indirectly; those who are the principals in the accomplishment of the valiant deed are better than those who are mere accomplices; those who carry the mission through are better than those who attempt but fail at it. What is more, even some fairly intricate "blaming" rules have their mirror image on the praise side. Take for instance the so-called "legal impossibility" doctrine, which says that a would-be criminal who commits what he mistakenly believes is a crime cannot be found guilty of a criminal attempt, because the crime he attempted is "legally impossible." By the same token, the would-be criminal's mirror image, the would-be hero who attempts what he mistakenly believes is a valiant deed, will not thereby earn much glory.

All of this should not be taken to imply that the symmetry between praise and blame is perfect. There are in fact many intriguing asymmetries. Here are just a few of the more noteworthy ones. Take first the defense of insanity. If the insanity defense operated symmetrically as between blame and praise, we would have to think worse of Van Gogh's masterpieces, once we realize that if he had committed crimes in the same frame of mind in which he "committed" his pictures, he could not be convicted by reason of his insanity. In fact, however, if anything, his insanity seems to increase the admiration we feel for his accomplishment. Consider next the role of risk and effort in the blaming and in the praising context. If they operated symmetrically, then just as we think better of the person who braved many risks and expended much effort to achieve his worthy goals, we would have to think worse of the criminal who braved many risks and expended much effort to achieve his unworthy goals. Which is, of course, the opposite of what we do. Consider lastly how luck and unintended consequences seem to play a different role with respect to praise than they do with respect to blame. Just think of Newton's

famous dictum that if he had seen farther than others, it was because he had stood on the shoulders of giants. It is indeed a well-accepted principle of praise that, as the mathematician Jacques Hadamard put it, "[w]hen the discoverer of a certain fact hears that another scholar has found a notable consequence of it, if this improvement has required some effort, the former will consider it not a failure but a success: he has the right to claim his part in the new discovery."⁴ This, however, would not seem to hold true on the blame side: We would not generally blame a person who has committed a bad deed, which then has further bad repercussions, for those repercussions.

Although the symmetry between praise and blame is thus far from perfect, it is extensive enough that we can bring a good deal of order to our previously chaotic intuitions about positive desert and answer a fair number of questions about desert that previously seemed unanswerable. Questions such as whether Copernicus's credit is diminished by the fact that in light of the physics he knew, his unswerving belief in his own system was probably unreasonable; or whether Werner Heisenberg deserves credit for denying Hitler the atom bomb by declining to invent it for him;⁵ or whether Frederick Banting deserved the Nobel Prize in Medicine despite the peculiarly circuitous route by which he stumbled onto the discovery of insulin.⁶ Being able to answer those questions makes the notion of desert much less vacuous. It means that desert has become commensurable: We are able to judge some rewards as being excessive or as being inadequate where previously we were "incommensurably" undecided.⁷

C. *Valuing Lives*

There is a great deal of inconsistency in the way we value lives. Some scholars have tried to account for this as a manifestation of incommensurability. Thus Cass Sunstein writes in his article on law and incommensurability:

There are extraordinary disparities in federal expenditures for each life saved. Some environmental programs prevent risks at enormous cost; the government is willing to spend relatively little to stop other

⁴ JACQUES HADAMARD, *THE PSYCHOLOGY OF INVENTION IN THE MATHEMATICAL FIELD* 49 (1945).

⁵ See THOMAS POWERS, *HEISENBERG'S WAR* 151 (1993).

⁶ See ALEXANDER KOHN, *FORTUNE OR FAILURE: MISSED OPPORTUNITIES AND CHANCE DISCOVERIES IN SCIENCE* 156 (1989).

⁷ See LEO KATZ, *ILL-GOTTEN GAINS* 197-263 (1996) (discussing the measurement of rewards).

risks. All current efforts to produce uniformity in expenditures—understood, for example, as equivalent amounts of dollar expenditures for each statistical life saved—have failed.

Such phenomena may in the end reflect irrationality, confusion, poor framing of relevant questions, interest-group power, or sheer chance. But it would be useful to explore other possible explanations. When people are thinking in these various ways, exactly what are they doing?

We might hypothesize that [inconsistent] social valuation of environmental goods comes partly from an insistence that diverse social goods . . . ought to be valued in different ways. With this hypothesis, some apparent anomalies dissolve or become more readily explicable. Some people, for example, insistently rebel against the idea that we should see all of the following, environmentally related consequences as “costs”: unemployment, higher prices, greater poverty, dirtier air, more cancer, respiratory problems, the loss of species. If we understand all these things as “costs,” to be assessed via the same metric, we will disable ourselves from making important distinctions. It might be hypothesized that when people refuse to trade off environmental quality and other goods, they are making a claim about the diversity of goods and incommensurability. They are claiming that one set of goods is superior to another not in the sense that it is infinitely valuable, but in the sense that it stands in a hierarchy of public values.⁸

Let us examine more closely what Sunstein has in mind with the help of a famous example from Guido Calabresi:

After about a month of studying cases, I put to my first term torts students a couple of hypothetical questions. The first concerns an “evil deity.” “Suppose,” I ask my students, “such a deity were to appear to you, as president of this country or as controller of our legal system, and offer a gift, a boon, which would make life more pleasant, more enjoyable than it is today. The gift can be anything you want—be as idealistic, or as obscene, or as greedy as you wish—except that it cannot save lives.” Later I will drop even that requirement. “The evil deity suggests that he can deliver this gift in exchange for one thing . . . the lives of one thousand young men and women picked by him at random who will each year die horrible deaths.”

When I ask; “Would you accept?” my students almost uniformly answer, “no.” Indeed, they are shocked that one could even ask the question. I then ask, quietly, what the difference is between this gift and the automobile, which takes some fifty-five thousand lives each year.⁹

⁸ Cass R. Sunstein, *Incommensurability and Valuation in Law*, 92 MICH. L. REV. 779, 835-37 (1994) (footnotes omitted).

⁹ GUIDO CALABRESI, *IDEALS, BELIEFS, ATTITUDES, AND THE LAW* 1 (1985) (ellipsis in original) (footnote omitted).

This is a miniature version of the kind of inconsistency Sunstein was talking about. Sunstein would like to account for this inconsistency by invoking incommensurability. He would suggest that a different metric applies when we are deciding whether to drive cars and when we are deciding whether to take the evil deity up on his offer.

At first glance, it is not clear how the applicability of different metrics brings incommensurability into the picture. In fact, it is not clear that incommensurability lurks anywhere in Calabresi's example. To be sure, the students who abhor the evil deity's offer, but who have no trouble embracing the automobile, do indeed seem to be applying a different "metric" when they are thinking about lives lost at the deity's altar and lives lost on the road. The two kinds of losses are not judged to be equal: The one kind of loss is considered worse than the other. That is why one is tolerable and the other is not. By making that judgment, however, the students show the two losses to be quite commensurable. So how is it that Sunstein manages to find incommensurability in such a situation?

Incommensurability gets into the picture once one takes account of the hesitation students feel when Calabresi confronts them with his hypothetical, and after they reject the evil deity's offer but welcome the automobile. Then they start to reflect on the seeming oddity of their position. They start to find it strange that it should matter whether a loss is inflicted by having everyone drive a car or by having the deity randomly pluck his victims. The students also might start to notice that most people would prefer to live in a society in which an evil deity plucks a thousand people at random each year than in a society in which accidents kill 55,000 each year. They might further think about the fact that if by spending a certain amount of money they could either eliminate all traffic accidents or eliminate the evil deity's annual prey, they would probably choose to spend it on traffic accidents. Once they think about these things, the students will start to feel undecided about whether it would be worse to accept the evil deity's offer than to allow cars to be driven. And that indecision has all the earmarks of incommensurability. The students have no trouble judging that an offer by the deity exacting only 999 lives would be better than one exacting 1000 lives. But that does not mean that they therefore find the deity's 999-lives offer preferable to an arrangement whereby 55,000 people die in traffic accidents each year.

My claim is that this incommensurability, like all the previous ones I have discussed, can be traced to ignorance and can be eliminated by dispelling that ignorance. What the students are ignorant of

are the properties of a nonconsequentialist, deontological morality. Once they become familiar with these properties, they will no longer view the choice between the deity's offer and the automobile as an incommensurable one, but happily will follow their initial intuitions—to reject the deity and choose the car. To a deontologist, there is nothing peculiar about treating the two kinds of arrangements differently. To each of the worries that made the students hesitate, the deontologist has an answer.

Start with the students' worry that it is strange that it should matter whether a loss is inflicted by having everyone drive a car or by having the deity randomly pluck his victims. To a deontologist, the way in which a loss is inflicted is hugely important. The hallmark of the deontological creed is that certain kinds of tradeoffs are off-limits. You cannot cut up one person to use his organs to save five others. You cannot execute one prisoner to prevent the mob from rioting and in the course of its riot killing several more innocents. You cannot torture the terrorist's child to get him to reveal where he has hidden the bomb. And so on. But when the deontologist prohibits carving up the one for the sake of the many, or killing a prisoner to protect innocents, or torturing a child to discover the whereabouts of the bomb, he acknowledges that different ways of bringing about a harm are morally different. When I do not carve someone up, I cause the death of five. When I do carve him up, I cause the death of one. Presumably *the way* in which I cause the death of the five (by letting them die) is, by the deontologist's lights, better than the way in which I would be causing the death of the one if I decided to save the five (namely by plunging a knife into him). This contrast, then, makes it no longer peculiar that we feel differently about deaths caused by allowing automobiles to be driven than by authorizing the evil deity to pluck a thousand victims each year.

Consider now the students' second reason for worrying whether accepting the evil deity's offer truly is worse than putting up with automobiles. They think about the fact that most people would prefer to live in a society in which an evil deity plucks a thousand people at random each year than in a society in which accidents kill 55,000 each year. This too is not very troubling to a deontologist. Most people would prefer to live in a world in which one can be carved up to save many, in which an innocent can be executed to appease a lynch mob, in which a child can be tortured to find the bomb. In such a society, lives saved are maximized, and by choosing to live in such a society you maximize your chance of being among

the survivors. This is simply a well-known counterintuitive aspect of the deontological position.

Consider, finally, the students' third source of worry. They think about the fact that if by spending a certain amount of money they could either eliminate all traffic accidents or eliminate the evil deity's annual prey, they probably would choose to spend the money on the elimination of traffic accidents. Does *that* not clearly indicate that putting up with automobiles is no better than accepting the evil deity's offer? The deontologist would point out that if 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 only stop 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.¹⁰

Once he has been able to make these arguments, the deontologist would, I think, have rendered the previously incommensurable choice between allowing society to use cars and accepting the deity's offer perfectly commensurable.

In the above three illustrations, I strove to show how one can trace cases of incommensurability to simple ignorance. It is worth pointing out, however, that one conversely can view many cases of ignorance as being instances of incommensurability. People will profess ignorance on many major moral questions: the death penalty, abortion, euthanasia, affirmative action, the use of lethal devices to protect private property, the felony-murder doctrine, and so on. Each of these can be thought of as a choice among incommensurables. We cannot decide between acceptance and condemnation, but we are not in equipoise between them either. How do I know we are not in equipoise? Well, just consider for a moment the person who cannot decide whether to accept or condemn the death penalty. He surely would find it easier to accept the death penalty for mass murderers than for simple killers. Yet he probably would continue to feel undecided about whether to accept or condemn the death penalty even for mass murderers.

¹⁰ For more on the three counterintuitive attributes of a deontological morality, see KATZ, *supra* note 7.

II. HOW TO JUDGE DECISIONS MADE IN A STATE OF UNAVOIDABLE IGNORANCE

When incommensurability is traceable to ignorance, it will often of course be but a temporary state—with knowledge, it will disappear. The problem is that decisions may have to be made before knowledge arrives. What then? How is one to decide the matter? The answer would seem to be: the best one can. And, if one truly has nothing to go on, then one might as well throw dice (as Donald Regan has suggested¹¹) or even resort to Mae West's solution. "Given the choice between two evils," she said, "I choose the one I haven't tried before." But what happens when knowledge does at last arrive and the choice that we have made in its absence turns out to have been the wrong one? The answer seems to be: nothing. If the decisionmaker did the best he could, it seems he cannot be blamed. Or can he?

In the remainder of this Article, I shall try to persuade you that very often the decisionmaker can be blamed all the same. The cases of incommensurability that have most interested philosophers are cases of moral, as opposed to factual, incommensurability. Although cases of factual incommensurability exist aplenty—as I pointed out in Part I, every litigated case, prior to the presentation of evidence, is an instance of factual incommensurability—it is cases of moral incommensurability that truly exercise people, like the questions of desert and the valuation of lives discussed earlier. If I am right that those are basically instances of moral ignorance, the person who makes what proves with hindsight to have been the wrong choice between "moral" incommensurables, will have to argue that his moral ignorance should excuse him. There is no doubt that the person who makes the wrong choice by reason of factual ignorance will be excused. But moral ignorance is a different story. If I shoot you because I did not see you standing behind the curtain, I am free of blame. I have acted out of factual ignorance. But if I shoot you because I did not realize that I may not use lethal force to defend my property, I have acted out of moral ignorance, and that is a quite different matter.

The question then is whether one can be blamed for a moral mistake that one could not help making. Strange though it seems, the criminal law has answered this question in the affirmative—that is what it means to say that ignorance of the law is no excuse. But does

¹¹ See Donald Regan, *Value, Comparability, and Choice*, in INCOMMENSURABILITY, INCOMPARABILITY, AND PRACTICAL REASON, *supra* note 2, at 129.

anyone really believe that principle, especially when it is construed as a moral principle, as saying that ignorance of a moral law—or just simply ignorance of morality—is no excuse? Is the principle not simply a crude utilitarian means for discouraging people from remaining willfully ignorant of the law so as to get away with their misdeeds?

No, it really does seem as though the principle comports with the way we judge misconduct. Think back to the infamous Milgram experiment.¹² Most of the people who obliged the experimenter and continued pressing buttons that they thought would inflict severe electric shocks on the uncooperative learner—those concentration camp guards *manqués*—did not enjoy doing what they did. In fact, when the experimenter gave them half a chance to cheat and not inflict the required shocks, they did. They went through agonies. They only continued pressing the buttons because they felt that that was what morality required of them. Yet, we are willing to condemn them for their weakness, despite the fact that it seems to have been the product of moral ignorance. To be sure, some might argue that we blame them only because they *should* have known better; that is, we blame them for being culpably ignorant of the moral thing to do here. So this is not a true case of deep moral ignorance. But my sense is that we do not blame them merely for being culpably ignorant; we blame them for doing the wrong thing, period.

For that matter, think about the way we evaluate the truly wicked, the likes of Stalin or Mao. Each of them surely thought he was acting morally, yet we have no trouble condemning them. And surely we condemn them not merely for having been culpably ignorant of what morality required of them. We also condemn them for their egregious immorality, without the least bit of concern for their indisputably sincere belief in their own moral rectitude.¹³

The idea that moral ignorance does not excuse gets indirect support from the way we think about judicial decisionmaking. Consider how the task of judging would change if we decided to treat ignorance in matters of legal doctrine exactly like ignorance of fact:

1. Suppose a judge does not know the answer to a legal question. He will make it up. He will reach for what seems to be the best solution even if it has low odds of being the right one. Not so with factual

¹² See STANLEY MILGRAM, *OBEDIENCE TO AUTHORITY* (1974).

¹³ Here, too, my sense is that we do not blame them merely for being culpably ignorant, but also for doing the wrong thing, period.

matters. If he cannot make up his mind about the facts, he is obliged to let the defendant win, to dismiss the case for lack of evidence. But a dearth of compelling legal arguments is not grounds for letting the defendant win. If we treated moral and factual ignorance alike, it would be.

2. The doctrine of precedent would disintegrate if the judge treated arguments (and their absence) like evidence (and the absence thereof). He might not be able to decide the case based on what he had heard today, but if convincing arguments were marshaled in the same kind of case tomorrow, he would be free to decide such a case differently then.

3. The doctrine of vagueness would become unnecessary. If a judge cannot decide whether a vague-looking statute covers what the defendant has done, he would simply say that in that case he has not been convinced beyond a reasonable doubt (if it is a criminal case) that the defendant is covered by the statute, and acquit. He would continue to apply the statute in cases in which he has no (reasonable) doubt that it applies. Even a statute as vague as one that simply made it a crime to engage in "terrible conduct" would not be declared unconstitutional. Its application would simply be confined to cases in which the judge was sure beyond a reasonable doubt that the defendant had acted terribly.

4. A whole range of subtler, more intricate consequences would follow as well. For instance, imagine the case in which exculpatory evidence appears after the defendant already has been convicted in a criminal case. Exculpatory evidence ordinarily is grounds for reopening a case. New exculpatory legal arguments are not, yet they would have to be if we put questions of law and fact on a par.

5. Prosecutorial discretion would be significantly narrowed. Currently, a prosecutor who wants to argue for a new legal theory is free to do so. If he presses a case on flimsy evidence, however, he is abusing his discretion. If moral and factual mistakes came to be treated alike, that asymmetry would have to be abolished.¹⁴

Despite everything I have said, I trust that the idea that factual and moral ignorance really are not very different probably will con-

¹⁴ For a superb exploration of the related issue of burdens of proof as they relate to issues of law and fact, see Gary Lawson, *Proving the Law*, 86 NW. U. L. REV. 859 (1992), and the fascinating replies it elicited from Larry Alexander and Richard Friedman, see Larry Alexander, *Proving the Law: Not Proven*, 86 NW. U. L. REV. 905 (1992); Richard D. Friedman, *Standards of Persuasion and the Distinction Between Fact and Law*, 86 NW. U. L. REV. 916 (1992).

tinue to have a very strong intuitive appeal. There is something profoundly strange about claiming that the person who acts wrongly out of unavoidable moral ignorance is to be blamed for his actions. If the actor has consulted every philosopher and judge on an issue, followed their advice, or gotten the advice that no one knows what the answer is, how can we possibly blame him for getting it wrong? To blame him seems to many downright irrational.

In fact, it is worth noting that one of the most famous Enlightenment arguments for religious tolerance appeals to the intuitive irrationality of treating ignorance in religious, moral, or legal matters any differently from ignorance in factual matters. The argument comes in the form of a parable first invented by Giovanni Boccaccio and presented as one of the many stories in his thirteenth-century classic, *The Decameron*.¹⁵ It was later refined and adapted by the eighteenth-century German playwright Gotthold Ephraim Lessing, who put the argument into the mouth of the eponymous hero of the best known German play on religious tolerance, *Nathan the Wise*.¹⁶ The Jew Nathan (Melchizedek in Boccaccio's version) comes to "invent" the parable when he is approached by the Sultan, who wants to intimidate him into lending him some money on favorable terms. The Sultan's means of intimidation is to ask the Jew to please tell him, since he is so wise, which is the right religion and why. Seeking to side-step that question, the Jew tells the Sultan a story, which, stripped of all poetic ornament, is this: There once was a very precious ring which, in addition to being very precious, entitled its bearer to be looked upon as heir and head of the family. This ring had been passed through many generations when it came to rest in the hands of a man who could not make up his mind as to which of his three equally good sons should inherit the ring. He solved his dilemma by having a master craftsman make two perfect copies and handing each son a ring with the assurance that it was the real one. When he died, the three sons of course began to quarrel, each claiming to be the rightful heir and head of the family. Alas, concludes the Jew:

[F]inding that the rings were so alike that it was impossible to tell them apart, the question of which of the sons was the true and rightful heir remained in abeyance, and has never been settled. "And I say to you, my lord, that the same applies to the three laws which God the Father

¹⁵ GIOVANNI BOCCACCIO, *THE DECAMERON* 44 (G.H. McWilliam trans., Penguin Books, 2d ed. 1995) (n.d.).

¹⁶ GOTTHOLD EPHRAIM LESSING, *NATHAN THE WISE* act 3, sc. 7 (Bayard Quincy Morgan trans., Frederick Ungar Publishing Co. 1955) (1799).

granted to His three peoples, and which formed the subject of your inquiry. Each of them considers itself the legitimate heir to His estate, each believes it possesses His one true law and observes His commandments. But as with the rings, the question as to which of them is right remains in abeyance."¹⁷

At first glance, there seems to be nothing particularly powerful about this parable. There seems no very compelling parallel to be drawn between the three religions and the three rings. Why should God be conceived to be akin to this irresolute father? Religious knowledge does not seem to be like a ring that, of necessity, only can be given to one person. Why did this story satisfy the Sultan? Why is the telling of it considered such an ingenious response to the Sultan's query? The ingenuity lies, I think, in the way the story appeals to our intuitive sense that factual and moral (and doctrinal and religious) ignorance are on a par. When we cannot settle a factual dispute, we normally will not be at each other's throats. Instead, we recognize that there is not enough evidence to determine the truth, and invoke the rules that are meant to deal with insufficient evidence. Of course, people often will go for each other's throats when a moral, religious, or doctrinal matter is at stake. The parable of the rings managed to convert the moral, religious, doctrinal question of which is the right religion into the factual question of which is the right ring. And on the ring question, everyone is willing to be tolerant in the absence of sufficient evidence.

The intuitive appeal of equating moral and factual ignorance thus cannot be denied. To be sure, I have tried to show you that a large number of quite strange implications are entailed by that position, but I fear they have not been nearly enough to overwhelm the strong countervailing intuition. Hence, the next Part of this Article will focus on that task.

III. CAN WE REALLY EVER BLAME THE UNAVOIDABLY IGNORANT?

If, out of genuinely unavoidable ignorance, someone makes what, in retrospect, turns out to be the wrong judgment call, blaming him seems like a cruel joke. It seems like holding someone responsible for murder who was tossed involuntarily on top of another person, thereby crushing him. To show why it is not necessarily unjust to blame someone who acted out of unavoidable moral ignorance, I will

¹⁷ BOCCACCIO, *supra* note 15, at 44.

reveal this to be but one manifestation of a more general, highly counterintuitive, aspect of deontological morality.

To get at this more general aspect, consider the following schematic scenario. Person One and Person Two live in parallel universes, so to speak. They behave identically in their parallel universes. Then each comes to a fork in the road. Indeed, they come to the very same fork in the road. At this fork, they can make a wrong decision or a right decision. Person One does the wrong thing. Person Two does the right thing. Which of them is worse? The answer seems to be self-evident. Indeed, it seems a merely terminological issue—an analytical truth. It seems clear that Person Two is better than Person One in the same way that we know it to be true that a bachelor is not married. Yet, that turns out to be false.

I begin with two simple examples that seem to contradict the common wisdom. There is much to quarrel about in each of these examples. Many reasons will suggest why the examples do not in fact show what they appear to show. But I will not linger to respond to these doubts here.¹⁸ I do not mean for those examples to *prove* the common wisdom wrong. I only mean for them to make it *conceivable* that the common wisdom is wrong. Having gotten you to take that possibility seriously, I then proceed to an actual proof—a proof that in a deontological system of morality, Person One, who has done more wrong than Person Two, might nonetheless end up with a better moral ledger!

My first example derives from William Manchester's biography of Douglas MacArthur, which describes a certain curious incident early in MacArthur's career.¹⁹ MacArthur was asked to serve in a noncombatant staff position, in which it was felt that his intellect could make the greatest contribution to the war effort. He politely, but resolutely, declined. MacArthur wanted to serve in the field. That was where glory lay; that was where medals could be won. Consider this incident just a bit more closely. If we believe, as I do, that the bestowal of medals pretty well tracks moral achievement, we would have to say that the valorous MacArthur is morally superior to the MacArthur

¹⁸ I presently am working on a piece in which I take up each of these possible exceptions in painstaking detail. See Leo Katz, *Preempting Oneself: The Moral Problems of Self-Manipulation* (unpublished manuscript, on file with author).

¹⁹ See WILLIAM MANCHESTER, *AMERICAN CAESAR: DOUGLAS MACARTHUR* 91 (1978). I am assuming that it is predictable from the very outset that he will do more good on staff but will achieve more glory in the battlefield. There is no element of fortuity involved here.

who decides to stay behind and serve on the staff, even though MacArthur actually might have saved more lives by staying behind, albeit at little risk to himself, and hence, with little claim to glory. Yet one also has the intuition that if MacArthur really could have accomplished more on staff than in the battlefield, that is where he should have stayed. If those two intuitions are correct—and at this point many will simply think one of them to be ill-considered—we would then have an example of the desired sort. The real MacArthur is Person One. The hypothetical MacArthur is Person Two. The real MacArthur makes the wrong decision. The hypothetical MacArthur makes the right decision. Yet, in the end, the real MacArthur comes out as morally superior to the hypothetical MacArthur.

My second example initially seems extraordinarily contrived. It will seem less contrived if I tell you what inspired it: the familiar problem under the necessity defense as to what we should do if the person claiming the necessity defense is in some way culpable in bringing such a predicament about in the first place.

Person One and Person Two both have poisoned five people. The poison does its work in a curious, but analytically helpful, way. The poison wrecks the hearts of some patients, the lungs of others, the kidneys of others yet. As things stand, each of the two sets of five victims contains one person in need of a heart transplant, two persons in need of lung transplants, and two persons in need of kidney transplants. At this point, Person One and Person Two have pangs of conscience. They decide they would like to save their victims. Both Person One and Person Two now face the familiar predicament of whether it would be all right to carve up an innocent to save the five. The answer for a deontologist is the familiar one as well: Of course not. This then is the fork in the road. Should one cut up one person for the sake of the many? And one branch of the fork is clearly right; the other is clearly wrong.

Let us suppose that Person Two decides to do the right thing. He does not carve up the one for the sake of the many. Therefore, his five victims die. He is now guilty of five murders. Person One decides to do the wrong thing. He carves up an innocent person and saves his five victims. He is now guilty of only one murder and several attempted murders, presumably a somewhat better ledger. Lest you have some doubts whether one murder plus five attempts is not as bad as five murders, we could modify the example slightly. Assume that Person One only tries to kill four people, two of whom are going to die unless they get kidney transplants and two of whom are going

to die unless they get lung transplants. He then assaults two people and extracts from each a redundant lung and kidney. He is now guilty of two assaults and four attempted murders. His counterpart, Person Two, by contrast, is guilty of four murders. Pretty clearly, two assaults and four attempted murders are less grave than four murders.

To repeat: I acknowledge that these examples invite all sorts of objections, which I will not here pass by. Instead, I will pass on to my proof that this sort of thing is an inevitable feature of deontological moral reasoning.

Consider Person One who commits Bad Act *A*. After he has committed Bad Act *A*, he commits Good Acts *B*, *C*, *D*, *E*, and *F*. Compare him to Person Two who has committed neither Bad Act *A*, nor Good Acts *B*, *C*, *D*, *E*, and *F*. Which is better? That depends. But one can certainly imagine Good Acts and Bad Acts such that Person One would come out ahead of Person Two despite the fact that he started out with a Bad Act, namely if the Good Acts are sufficiently good, sufficiently courageous, and beneficial. Assume we are dealing with such an *A*, *B*, *C*, *D*, *E*, and *F*.

Consider now Person Three. Person Three faces a situation in which he has a choice either to do nothing—commit neither Bad Act *A*, nor Good Acts *B*, *C*, *D*, *E*, and *F*—or to do both *A*, and *B*, *C*, *D*, *E*, and *F*. He decides to do *A* so that he can also do *B* through *F*. How does he compare with Person One? He did the same as Person One, except that he did *A* so that he could do *B* through *F*. So, he seems a little better, but certainly no worse than Person One. He thus is clearly better than Person Two.

But now consider more closely his decision to do *A* so as to be able to do *B* through *F*. Is that a correct decision? Under a deontological regime, it generally will not be. Let *A* = torturing a child, and *B* = a valiant act of rescue of several people, *C* = the same thing, but involving other people, *D* through *F* = ditto. Torturing a child so as to rescue several people is impermissible to a deontologist. (To be sure, most people are threshold deontologists, and if the number of people to be rescued is sufficiently large, then the tradeoff is permissible. I am assuming that we are below that threshold.) Now it seems clear that Person Three is making a decision that is morally inferior to the decision made by Person Two. Yet it is also clear that he comes out ahead morally.²⁰

²⁰ Notice that nothing about this paradox depends on the so-called "moral luck" phenomenon. When first encountering cases of the kind I discuss in this Part, espe-

This example suggests a way in which we can make sense of our seemingly conflicting intuitions about moral ignorance. We could regard the person who acts on the best available moral advice, but still gets it wrong, in the same light we regard the protagonists of my various examples. The morality of what he does at the moment he acts is to be viewed like the morality of the person who declines to cut up one to save five (whom he tried to poison initially), or of the hypothetical MacArthur who decides to serve on staff rather than in combat, or of the hypothetical variant of Person Three who decides he is not going to commit Bad Act A just to be able to then commit Good Acts B, C, D, E, and F. He, like they, could be thought to have acted rightly but, nonetheless, by that self-same right action, to have *blackened* his moral ledger. On the face of it, that is an absurd suggestion. What my examples are meant to show is that far from being absurd, it is an almost inevitable feature of a deontological system of morality.

CONCLUSION

I have argued that incommensurability often is simply a sign of ignorance, most frequently ignorance of the answer to difficult moral questions. When one has to choose between incommensurable alternatives, what one is really doing, most of the time, is attempting to answer a hard moral question. An interesting and unnoticed aspect of this choice is that if one gets it wrong, and if, with the passage of time and the achievement of greater moral insight, that becomes

cially when encountering my surgery hypothetical, it is tempting to think that the paradox therein simply derives from the fact that we attach moral significance to chance consequences. It can be shown that this is in fact not true of the surgery example, and I go to great trouble to do that in my forthcoming piece mentioned above. See *supra* note 18. But most obviously it is not true of the Good Act-Bad Act hypothetical. No chance events are involved there; everything happens exactly as intended.

This is also a good place to correct two misconceptions that frequently obscure an understanding of deontological morality. The first misconception is that deontologists necessarily embrace moral luck, that the deontologist necessarily attaches moral significance to chance consequences. In fact, both kinds of deontologists exist—those who do, and those who do not, attach such significance. The second misconception is that a deontologist attaches no significance to consequences, that only the action component of misconduct is morally relevant. That is not true either. A deontologist is someone who attaches significance to things *in addition* to consequences, namely the way in which a consequence is brought about.

At least the above is true of the kind of deontological morality I am discussing in this Article. I believe that deontology of this stripe coincides both with our everyday moral judgments and with the underlying logic of our criminal law.

apparent, one still can be condemned for that choice. This condemnation is justified even though one really could not help making the mistake. We see here a surprising vindication of the commonly held intuition that incommensurability and tragedy are deeply intertwined.

* * * * *