Theoretical Inquiries in Law 4.1 (2003)

Harm and Justification in Negligence

Leo Katz[*]

Negligence, the creation of an unjustifiable risk of harm, plays a pivotal role in both criminal and civil law. This article takes up two negligence-related problems unique to its role in the criminal law. The first has to do with its "harm" component, the second with its "unjustifiability" component. The first problem is why the criminal law distinguishes so sharply between negligent wrongdoing that results in harm and negligent wrongdoing that does not, when it does not distinguish equally sharply between intentional wrongdoing that results in harm and intentional wrongdoing that does not. The answer has to do with a peculiar feature in the logic of intention and with the symmetric way we ascribe responsibility for blameworthy and praiseworthy actions. My second problem is why the criminal law judges the blameworthiness of negligent actions on the basis of the harm wrought rather than the gap between the harm wrought and the justification for risking it. The answer has to do with a more general feature of the criminal law, the refusal to recognize "partial defenses" and "partial offenses."

I. TWO PROBLEMS

Negligence figures in both torts and criminal law; and the concept looks quite similar in the two contexts. But there are some important problems associated with negligence that are unique to criminal law, and it is two of those that I will take up here.

---

* I want to thank my fellow symposiasts for their invaluable reactions, Shai Lavi for his astute and painstaking comment, and Claire Finkelstein for innumerable suggestions and improvements.

My two problems arise out of the fact that criminal law, unlike tort law, has to concern itself with blameworthiness. Criminal law is about retribution, and retribution requires punishment to be proportional to the blameworthiness of the defendant's actions. That means that the criminal law has to answer a set of questions the tort law does not have to contend with. In tort law, we only need to find out whether the defendant acted wrongfully and what compensable harm that wrongful conduct caused. In criminal law, we also need to find out just *how bad* that wrongful conduct was. In tort law, it does not really matter what kind of wrong someone committed. (As long as it constitutes a tort, be it minor or major, the wrongdoer must pay for its harmful consequences.) In criminal law, the nature of the wrong, the seriousness of the crime, is of the essence. It determines the proper measure of retribution. When the wrong happens to be negligence, the criminal law, unlike the tort law, must, therefore, determine not merely *whether* the defendant is guilty of negligence, but also *how bad* his negligence was. My two problems both involve the making of the latter determination — how bad the defendant's negligence was.

The Model Penal Code defines negligence as the imposition of a substantial, unjustifiable risk.[1] This definition is basically identical to that used in tort law, although it employs somewhat different language than is customary in the latter. Tort law defines negligence as an act whose probability of harm exceeds its benefit, but these are, of course, just slightly different terms for risk and justifiability.[2] Negligence is thus a function of just two things: risk and justifiability. It requires that there be enough of the former and too little of the latter. My first problem has to do with the role that the second of those components, justifiability, plays in assessing how bad someone's negligence has been.

Consider the case of someone who speeds his car through a busy market square in order to take an injured passenger to the emergency room. Suppose he runs into several pedestrians and ends up killing them. Is he guilty of manslaughter? That depends on the risk his driving posed (the likelihood that he was going to run several of them over) and on his reason for running it (the severity of the passenger's injury and the need to get him to a hospital quickly). If the defendant's reason for speeding is compelling

---

1 Model Penal Code § 2.02 (Proposed Official Draft 1962).

2 A technical aside: criminal law distinguishes between two kinds of negligence, knowing negligence, which the Model Penal Code calls recklessness, and unknowing negligence, which it calls negligence pure and simple, *id.* The Model Penal Code generally insists on knowing negligence as a prerequisite of liability, and it is the only kind of negligence I will be discussing.

enough to justify the risk he created, everything is simple: he is not guilty of manslaughter. But if it is not compelling enough, in other words, if it falls short of justifying his actions, things get more mysterious. For suppose that his reason falls *just barely* short of justifying his actions. Then he is guilty of manslaughter. Just like someone who has no justification at all for what he did. Maybe we regard him as a little better than someone who has no justification at all, but not very much better. We give the defendant virtually no credit for having a justification that falls *just barely* short of exonerating him. If we were to do so, we would find him no more blameworthy than a jaywalker, since his justification is so very nearly enough to justify him. The fact that we think him a whole lot worse than a jaywalker shows that we do not give much "partial credit" for an inadequate justification. And that seems strange and cries out for an explanation. I will call this the "partial credit" problem.

A slightly different way of putting the partial credit problem would be thus: We determine whether there was negligence by comparing the risk with the justification for it. However, once we have concluded that there was negligence, we then determine how bad that negligence was by focusing exclusively on the risk, rather than on the "gap" between the risk and the justification. Intuitively, it seems as though the "gap" would be a more natural measure of blameworthiness. After all when that gap is zero, we say that there is no negligence. Therefore, when that gap is slight, it seems we should say there was but slight negligence, and if it is substantial, we should say there was substantial negligence. But what we do instead is to say that once the gap is greater than zero, then the size of the risk, rather than the size of the gap, will be the determining factor in assessing blameworthiness (and punishment). What that means concretely is that if we compare a defendant who negligently created a risk of death to several people, but for a good, though not quite sufficiently good, purpose, and a defendant who negligently created risk to some valuable property, but for no good purpose whatsoever, the former is automatically worse than the latter: he created the bigger risk. Never mind that he had better reasons. We give no partial credit for insufficiently good justifications. Why is that? And should it be that way?

The partial credit problem relates to the justification component of negligence. My second problem relates to the risk component. Compare two negligently speeding drivers. The first runs someone over; the second does not. The first might be guilty of manslaughter; the second will be guilty at most of reckless endangerment, which tends to carry a mild punishment and should, in the eyes of many, not be criminalized at all. It is a well-recognized, as yet unresolved, mystery about such cases, why we should care whether

anyone was actually run over, as opposed to whether the driver merely risked running someone over. This is the issue of moral luck.[3] I will not have anything to say about it here. But there is another, less well-recognized mystery about such cases as well. Let us add two further drivers to the tableau. These two, unlike the previous two, actually want to run their victim over. As before, the first actually does so; the second misses. The first is guilty of murder; the second of attempted murder, a less serious but still plenty serious infraction. That raises the following question: Why is it that when the two drivers are merely negligent, it is only the first driver who is guilty of anything serious, and the second driver barely blameworthy at all? In other words, when it comes to intentional wrongdoing, it only matters somewhat whether the risk the defendant created actually materializes. Even if no harm comes to pass, he is still guilty of a serious crime: attempted murder. By contrast, when it comes to negligent wrongdoing, it matters greatly whether the risk actually materializes. If no harm occurs, he is barely guilty of anything at all. The absence of actual harm makes much more of a difference for negligent wrongdoing than for intentional wrongdoing. Why is that? I call this the "harm problem."[4]

Because the harm problem is so much more easily disposed of than the partial credit problem, I will deal with it first.

II. THE HARM PROBLEM

The familiar way of putting my first problem is to ask why the criminal law does not treat every uncompleted crime, every "crime-fragment," as it were, as a criminal attempt. If the defendant has exhibited the *mens rea* of a certain crime and has completed a substantial portion of its *actus reus*, though he was unable to complete it, why do we not judge him guilty of attempting that crime? To be guilty of an attempt requires an intent to commit a crime, and that means that a reckless actor who comes within a hair's breadth of completing a crime of recklessness cannot be found guilty of attempting

---

3 The two articles reawakening this problem to philosophical life in the last few decades are Thomas Nagel, *Moral Luck, in* Mortal Questions 24 (1979), and Bernard Williams, *Moral Luck, in* Moral Luck 20 (1981). For more recent treatments, see Michael S. Moore, *The Independent Moral Significance of Wrongdoing, in* Placing Blame 191(1997), and Leo Katz, *Why the Successful Assassin Is More Wicked Than the Unsuccessful One*, 88 Cal. L. Rev. 791 (2000).

4 For recent explorations of this question, see Anthony Duff, Criminal Attempts (1996).

the crime he just fell short of committing. This has struck many people as paradoxical.

For at least some cases, a sort of explanation quickly suggests itself. Take J.C. Smith's hypothetical about "a driver, X, [who] overtakes on the brow of a hill when he cannot see whether anything is coming in the opposite direction. If anything were coming, there would certainly be a crash and, obviously someone might be killed. But nothing is coming and no harm is done."[5] Smith finds it "startling" to contemplate convicting X of attempted murder. But why exactly? I suspect that what troubles Smith is the incredibly hypothetical nature of the harm we would be charging X with attempting. Once we are in the realm of hypothesizing a nonexistent car coming up from the opposite direction, why not hypothesize two, or three, such cars, and find X guilty of three attempted murders rather than merely one? For that matter, why not hypothesize a truck piled high with property of some kind, which a collision might well damage or destroy? Now the defendant would be guilty of some sort of attempted property offense. In other words, the trouble with charging unintentional actors with attempting unintentional crimes is that it will often be unclear which unintentional crimes and how many unintentional crimes they could be charged with. As a result of this unclarity, any charges we do press will seem very arbitrary. Intentional crimes do not present this difficulty. As to them, it is always quite clear which and how many crimes the defendant is attempting, to-wit, the crimes he is intending to commit.

The trouble with this explanation is that many reckless crimes do not present this difficulty. Not all of them are like Smith's driving case. Consider the old chestnut about the assassin who tries to kill the queen by planting a bomb in her coach. It is likely that the bomb will kill both queen and coachman. If the bomb is found in time, the assassin will be found guilty of attempting to kill the queen, but not of attempting to kill the coachman, since he intended her death, but not his. Yet there really is no issue of arbitrariness or hypothetical crimes that is getting in our way here. If the bomb had gone off, the assassin would be guilty of both the murder of the queen and of the coachman. Therefore the only thing we could possibly regard him as guilty of, if the bomb does not go off, is the attempted murder of the queen and the coachman. Why, then, is it that we decline to do so? Why are we only willing to find him guilty of an attempt on the life of the former and not the latter?

---

5 John C. Smith, *The Elements of Chance in Criminal Liability*, 1971 Crim. L. Rev. 63, 72-73, *cited in* Sanford H. Kadish & Stephen J. Schulhofer, Criminal Law and Its Processes: Cases and Materials 590 (6th ed. 1995).

Some people might take issue with my claim that the assassin would not be found guilty of attempting to murder the coachman. They would argue that the assassin was not merely risking death to the coachman, but intending it. It is true that the assassin probably regards the coachman's death with indifference or even regret. But that, they will argue, does not prove that he did not intend it. The person who kills his grandmother so that he can inherit her money will not then be allowed to claim that he did not really intend to kill her on the grounds that the actions he took to get at the money she had willed to him, namely, poisoning her, only had her death as an unfortunate byproduct, his intent being directed toward getting at her money, not toward killing her. Her death, everyone would insist, was no mere byproduct of his actions; rather, it was his means of getting at her money and therefore fully intended. Someone might say the same about the queen's assassin: blowing up the coach and everybody on or in it was no mere byproduct of killing the queen — it was the means toward that end. Therefore the queen's assassin intended the coachman's death just as much as the grandson intended the grandmother's death.

Well, maybe. Or maybe not. The line between means and byproducts is notoriously vexing, and the distinction between intentional and unintentional actions correspondingly uncertain. That uncertainty, however, cannot be used to turn every case in which someone creates a high probability of harm to a particular person into an intentional attempt on that person's life. So our puzzle remains. Why do we judge a case in which someone creates a high probability of harm (but does not actually harm anyone) so much more leniently than a case in which he was actually hoping to inflict such harm? Our difficulty in distinguishing those two cases only serves to heighten the puzzle. It is strange that two cases that are so hard to tell apart should receive such drastically different evaluations. So strange, that many people have argued that the law gets it wrong and should be changed: the two cases should be treated roughly the same.[6]

But I think the law here does, in fact, get things right and should not be changed. One way to see this is to focus on a familiar but somewhat neglected oddity about the concept of intention. Imagine a defendant who aims at a faraway target and fires. Assume, to begin with, that he has no desire to hit anyone and that his chances of hitting that faraway target are so low that no one would judge his behavior reckless. If, by some miracle, his bullet nevertheless manages to travel much farther than anyone had reason to expect, he would not, therefore, be guilty of anything (either in law or

6 Such as Model Penal Code section 5.01(1)(b) (Proposed Official Draft 1962).

in morality). Next suppose that when firing the bullet, he does, in fact, entertain the desperate hope that the bullet will reach and kill that faraway target. Now he would be guilty of murder. A defendant who brings about a hoped-for low probability event is guilty of intentionally causing it. To be sure, the probability cannot be too low. There are levels of improbability at which we are not prepared to hold him responsible, regardless of the fervor or intensity with which he hoped. But if the improbability of the outcome is not quite that great, he will be held responsible for it. This principle — this oddity — that someone is responsible for extremely unlikely, hoped-for consequences holds not merely for bad outcomes, but also for good ones. Imagine a soldier, fighting for a just cause, who, at great risk and with great effort, takes aim at an enemy much too far away for him to have a reasonable hope of hitting, but he hits that enemy all the same. We have no reluctance crediting him with his miraculous success. We hold someone responsible for good or for ill for any however-improbable result, *so long as he hoped that it would come about.*

Let us apply this principle to the case of the defendant aiming to kill the queen, but expecting to kill the coachman along the way. When he plants his bomb, he sincerely hopes to kill the queen, but equally sincerely hopes that the coachman will escape unscathed. But then if the coachman does escape unscathed, contrary to expectations, perhaps because the bomb was found in time or because miraculously it only killed the queen, then the defendant can be said to have intentionally spared him. To be sure, sparing him was an unlikely outcome of the assassination venture. But however unlikely, it was hoped-for. And that makes it intentional. We therefore must credit the assassin with sparing the coachman's life. If we credit him with sparing the coachman's life, that makes it difficult to hold him responsible for attempting to kill him.

Perhaps this seems a little too quick. So let me try to get to the same result by a slightly different argument. It relies, at bottom, on the same idea, but has a more familiarly doctrinal flavor to it. The defense of abandonment provides that a defendant who is guilty of a criminal attempt escapes liability if he voluntarily desists from his criminal designs. This might happen fairly late in the game: the bank robber might have broken into the bank, opened the vault, removed the money, and only then experienced a change of heart, after which he put the money back into the safe and left. There is no doubt that at the point at which he removed the money from the vault he was guilty of a criminal attempt. But according to the abandonment defense, his subsequent voluntary discontinuation of his plans purges him of all guilt. Now suppose that the queen's assassin had initially wanted to kill both the queen and the coachman. However, after planting the bomb and seeing it go off, he

changed his mind about the coachman and dragged the man's unconscious but not yet dead body out of the conflagration. He might then well be entitled to claim the abandonment defense: he voluntarily desisted from his original plan to kill the coachman. This is hardly uncontroversial. Some people think that there is a point at which a criminal attempt should be deemed "unabandonable." Obviously it is unabandonable once the crime has been completed. But it is often suggested that it should be unabandonable even earlier. Courts, however, have had a real problem finding a principled basis on which to designate such a point, and many have, therefore, concluded for the most part that an attempt remains abandonable up until it has actually been completed.[7]

Suppose next that the queen's assassin never actually wanted any harm to befall the coachman and, in order to protect him, moved the bomb to a spot where there was at least a remote chance that the explosion would spare him, even while it killed the queen. And suppose that the coachman in fact survives. Now does the queen's assassin deserve to be found guilty of the coachman's attempted murder? It would be hard to argue that he is covered by the abandonment defense, since he did not change course midway through his crime. Nevertheless, it seems clear we need to give him some kind of special break. After all, the reason he does not qualify for the abandonment defense is that he "abandoned" his criminal designs vis-à-vis the coachman from the very outset, by taking certain modest measures to spare his life. To be sure, those measures were, *ex ante*, extremely unlikely to be of much help, but that is irrelevant. Someone who does qualify for an abandonment defense will often only have just managed to pull the chestnuts out of the fire by some, *ex ante*, extremely unpromising last minute rescue measures. If we are going to acquit the queen's assassin where he manages to rescue the coachman after the bomb has gone off, it seems clear we also have to acquit him if he tries to spare him in advance.

Consider finally the case in which the queen's assassin does nothing to decrease the likelihood that the bomb will kill the coachman, except hope against hope that the coachman will be spared. He plants the bomb where it is most likely to kill the queen and says a prayer for the coachman. Miraculously the coachman is unhurt. In this latest variation, does the queen's

---

7 In a particularly striking case of this sort that arose in Germany, the defendant tried to strangle his victim and, taking him for dead, started to walk away. Then he heard the victim gasp, "I am still alive!" whereupon he turned around and, with something between a sneer and a sigh of resignation, grumbled, "So live!" and left. He successfully claimed abandonment. BGHSt 36, 224 (known as the *Ich lebe noch-Fall*).

assassin deserve to be found guilty of attempting to murder the coachman? Only if there is a meaningful difference between this case and the previous one, in which he moves the bomb to a location where it marginally improves the coachman's odds. I cannot see enough of a difference to warrant treating them differently. If we acquit in the previous case, we should acquit in this one.

What the defendant in this last scenario has engaged in is an act of extreme recklessness toward the coachman, which, by some extraordinary luck, did him no harm. Nevertheless, we concluded that the defendant does not deserve to be punished. If this is right, then we here have our explanation why there are no "reckless attempts" and why actual harm is such a crucial prerequisite for the serious punishment of excessive risk-taking.

III. THE "PARTIAL CREDIT" PROBLEM

Negligence requires a substantial, unjustifiable risk. But once we have determined that the defendant was in fact criminally negligent (or reckless), the extent of his justifiability drops out of the picture. Only the substantiality of the risk he created will guide us in judging how badly he behaved and how harshly he should be punished. As between two drivers, one of whom speeds through a busy market square only to see the panicked crowds scatter and the other of whom is trying to get a not so critically injured patient to the hospital, the law does not find much of a difference. Both might well be guilty of manslaughter. Not that the second driver's reason counts for nothing, but it counts for little. He gets no "partial credit" for it. At least that is what my previous example about a driver whose justification falls just short of exonerating him suggests.

This is not the only instance in which we throw away information, as it were, in the course of judging negligence. There are some others as well. Perhaps our reasons for discarding information in *those* cases might also apply in *this* one? Imagine one such case. The defendant is racing his car to the emergency room on account of a passenger who has suffered a relatively slight injury. Let us suppose, however, that the passenger is a violinist and the injury, although slight, might significantly affect her manual dexterity and thus endanger her virtuoso status. If that were to happen, we will assume her life would be ruined: she might well commit suicide. Do these special features of her case count? Ought they to be thrown into the balance in determining whether the risk the driver imposed on the rest of the traffic was justified?

My guess is they ought not to be. To test that out, and to see what

implications it has, let us flesh this example out a bit more. Suppose the violinist arrives at the hospital and the doctor begins to examine her. In the midst of his examination, another patient is wheeled in. The man has suffered an injury to his left leg. If the injury is not quickly taken care of, he is likely to loose the leg. The doctor interrupts his work on the violinist to turn his attention to the injured man. The violinist protests. The exasperated doctor exclaims that surely someone who is about to lose the use of one of his legs has a greater claim to immediate attention than someone who merely has to fear a slight potential decrease in manual dexterity. But the violinist persists. "You don't understand," she says, "my violin-playing is tremendously important to me. If I am unable, or even just somewhat less able, to play it, that will really take much of the joy out of my life. By contrast, look at fellow over there. I happen to know him. He is depressed, without talents, without interests, and without prospects of making a contribution to either his own or society's happiness. You would do more for humanity by making sure my manual dexterity is not decreased by an iota than by seeing to it that he retains the use of his leg." Would it be all right for the doctor to acquiesce in this argument, disregard the injured man, and continue ministering to the violinist? Presumably not. What I mean to show by this is that in deciding on the right course of action, we do not pay attention to the competing intensities of desire of the parties involved. What we take into account instead of their desires are their *claims*: we compare the strength of their claims to assistance, or forbearance, or whatever it is that they want of us. What we think should matter to the doctor in deciding which patient to attend to first is not who would suffer the greatest utility loss as a result of delayed treatment, but who has the greater claim to his assistance. And that would be the man who risks losing his leg rather than the violinist who only risks losing some manual dexterity, because each of us has more of a claim to preserving his ability to walk than his ability to play the violin like a virtuoso. (Correspondingly, we think that what should matter to the driver in deciding how fast to drive to the emergency room is not whether the bystanders he might run into or the violinist he is seeking to rescue is threatened with the greater utility loss, but who has a greater claim to his solicitude. And that would presumably be the bystanders, who might get run over and incapacitated, rather than the violinist.)

Admittedly, this is bound to sound a little strange. How can we say that the claim someone has to retain his ability to walk is greater than someone else's claim to retain his ability to play the violin, when the walker cares so much less intensely about his walking than the violinist cares about his violin playing? But that really is not as strange as it sounds. Imagine that someone were to try to steal my attaché case, which happens to contain

a just-completed manuscript of mine that I treasure beyond anything else, almost more than life itself, enough, in any event, that I would gladly risk life and limb to rescue it. Could I use deadly force to try to retrieve it? Almost certainly not. Property, however treasured, may not be defended with deadly force. But why is that? Why is it that I can defend an attack against my person with deadly force but not an attack on some much-valued property, like that manuscript? Presumably because my claim to bodily integrity is greater than my claim to the integrity of my property. This holds true despite the fact that I value certain property more dearly than bodily integrity. Claims need not match up with preferences.[8]

Now it is well to quickly acknowledge that this can lead to certain paradoxical-seeming results. Suppose, for instance, that just as the doctor is about to attend to the injured man's leg, the violinist discovers that in her concern over her hand, she overlooked a far more serious injury to another part of her body, her spinal chord, let us say. If not immediately attended to, she risks losing the use of all of her lower extremities. She points this out to the doctor, who immediately agrees that that gives her priority over the injured man. Now suppose that as the doctor is about to attend to the violinist's spinal chord injury, she tells him that she would like him to first take care of her hand. Losing manual dexterity and violin-playing ability, she explains, would be more terrible for her than losing the use of her lower extremities, unimportant as they are for violin-playing. What now? The following three contentions all seem, superficially at least, true:

(1) As between looking after the injured man's leg and taking care of the violinist's spinal chord injury, surely the doctor ought to do the latter, since the violinist's claim to maintaining the use of her lower extremities is greater than the man's claim to maintaining the use of his left leg.

(2) As between taking care of the violinist's spinal chord injury and attending to her hand, the doctor ought to forget about the spinal chord and attend to the hand, since the violinist would prefer it that way, it is her body and we like to respect autonomy, and doing so in no way would disadvantage the other patient.

(3) As between attending to the violinist's hand and looking after the injured man's leg, surely the doctor should look after the leg, since the man has

---

8 For a defense of related ideas, see Thomas Scanlon, *Preference and Urgency*, 7 J. Phil. 655, 659-60 (1975); Thomas Nagel, The View from Nowhere 171 n.1 (1986); Ronald Dworkin, *What Is Equality? Part 1: Equality of Welfare*, 10 Phil. & Pub. Aff. 185 (1981).

more of a claim to getting his leg repaired than the violinist has with respect to her hand.

Unfortunately, if we combine these three contentions, they put the doctor into an endless cycle. So what is he to do? One of the statements will have to be rejected. Assuming that we believe what I have said about interests, we should probably reject (2). In other words, we should probably accept that even if the violinist asks that the doctor treat her hands rather than her spinal chord injury, he should say to her, "You can either get your spinal chord treated or you have to wait your turn while I treat the man's leg." We should not allow the violinist to trade-in, as it were, her right to have her spinal chord treated for treatment of her hand. As I indicated, this amounts to nothing short of rejecting a Pareto-superior alternative (and therewith the Pareto principle): for the doctor to treat the violinist's hand rather than her spinal chord would improve the violinist's position without worsening the position of the injured man. Nonetheless, we are forced to consider such an arrangement off-limits.

Notwithstanding this strange implication, I think the "claims" argument is right. In evaluating the justifiability of someone's risky conduct, we should be concerned with the strength of his claims, not with his utility or his preferences. Crudely put, we should be objective rather than subjective. But that means disregarding a great deal that the defendant might have marshaled in his favor. It disables him from justifying his risky conduct by invoking his intense attachment to whatever it was his risky conduct was meant to protect or effectuate.[9]

The "claims" argument serves to shed some light on the partial credit problem. It tells us why the defendant does not get partial credit for certain things that one might think partially justify his misconduct. It tells us why the fact that the defendant greatly values the thing for the sake of which he ran the risk does not count for much. It tells us why we do not place *his* valuation on the benefits with which he defends the loss he risked: why we do not consider the violinist's violin-playing ability as precious as *she* does when we decide how much risk may be imposed on others to protect it.

But there is also a great deal the "claims" argument does not tell us. The partial credit problem arises because in determining how blameworthy the negligent actor is, we are not merely refusing to place *his* valuation on the benefits he sought to achieve, we are refusing to place any value on them whatsoever. Put differently, we not only fail to give the defendant the credit

---

9 For a fuller discussion of such cases, see Leo Katz, Ill-Gotten Gains: Evasion, Blackmail, Fraud and Kindred Puzzles of the Law pt. 2 (1996).

he thinks he is entitled to, we fail to give him any credit at all. The "claims" approach explains why we do not give the defendant the credit he thinks he is entitled to; it does not explain why we fail to give him any credit at all. For that we will have to look elsewhere.

IV. THE "PARTIAL DEFENSE" PROBLEM

The partial credit problem is but a special case of a much more general phenomenon. Another manifestation of that phenomenon is to be found in the area of criminal law defenses: the justifications and the excuses. Think about the case of a defendant who claims duress, or self-defense, or necessity as a defense to her crime. Suppose that although her claim is a colorable one, we ultimately decide it is not sufficient. In other words, she falls just short of mounting a valid defense. How guilty is this defendant? As guilty as if she had no shred of a defense whatsoever. Here, as in negligence, we refuse to give partial credit. We could have said: someone who *nearly* qualifies for a duress defense because she was subject to extreme, but not extreme enough, pressures to commit a crime gets "partial credit" for those pressures and will have her guilt, and punishment, reduced accordingly. Put differently, we could try to measure her blameworthiness by the gap between the gravity of the crime she committed and the strength of her defense. But we do not. If the defense falls short, we are mostly guided by the gravity of the offense.

To be sure, there are exceptions: The provocation defense seems to give defendants partial credit for certain kinds of pressures they may have been subject to. Diminished capacity seems to give them partial credit for partial insanity. Imperfect self-defense gives partial credit for some kinds of invalid self-defense claims. And Douglas Husak recently suggested in a deeply insightful article called *Partial Defenses*[10] that really all defenses *ought* to be available on a partial basis, though he acknowledges that at present they are not. We might go even further than he suggests. Once one has embraced the partial credit idea, it is not clear why one might not conclude that someone who, let us say, has "half" of a self-defense justification and "half" of a duress excuse therefore has accumulated enough credit for a complete defense and should be acquitted. But I suspect that even Husak would be reluctant to recommend such aggregation of partial credits.

That should make the partial credit problem in negligence considerably less mysterious, because it is less *sui generis*. But it still leaves us with a

---

10 Douglas Husak, *Partial Defenses*, 11 Can. J.L. & Jurisprudence 167 (1998).

question. What is at the base of this more general phenomenon? Let us see if we can shed some light on this by looking at a few defenses more closely.

A. Duress

The Model Penal Code grants the duress defense to someone who commits a crime "because he was coerced to do so by the use of, or threat to use, unlawful force against his person or the person of another, which a person of reasonable firmness in his situation would have been unable to resist."[11] Why does it not substantially mitigate the defendant's guilt that he yielded to pressure that was almost but not quite so strong that "a person of reasonable firmness ... would have been unable to resist"? Would not a more reasonable approach be to give the defendant more and more of a break the more severe the pressure to which he was subjected, up until the point at which the pressure is so great that not even a reasonable person could withstand it, which is when we are ready to exonerate him altogether? Why is the law so thoroughly uncompromising — in the most literal sense of that term?

Let me suggest a reason. Consider the case of a defendant who is faced with the threat of significant property damage unless he commits perjury. Most of us would probably not regard that as a serious enough threat to trigger the duress defense. A threat against the defendant's life would, of course, be a different matter; but the threat to do, say, $10,000 of damage to his car does not seem to qualify for the defense. Should we nonetheless show the defendant some indulgence on the grounds that he is better than someone who perjures himself without having faced such a threat? It may seem obvious that he is, but there are reasons for thinking he is not. Imagine someone who has just lost $10,000 in gambling. Suppose that person were offered $10,000 in exchange for committing perjury. How would we feel about such a person? Ten-thousand dollars is admittedly tempting, especially for someone who has just lost that amount at the gambling table. But I do not think that would induce us to judge him any less harshly than had he committed the perjury gratuitously, as it were, without the benefit of such a payoff. Great temptation does not seem a good enough ground for diminution of guilt. Now compare the defendant who is threatened with the loss of $10,000 unless he commits perjury with the defendant who is offered the chance to make back his $10,000 gambling loss by committing perjury. Are they markedly different? A difference is not easy to make out. What that tells us is that the reason we decline to give a defendant "partial credit"

---

11 Model Penal Code § 2.09.

for a crime committed under serious pressure is that we would then have to give similarly "partial credit" for a crime committed under temptation. And that seems unpalatable.

You might wonder whether my argument does not prove too much. Is it not possible to apply the same argument to valid cases of duress and thereby "prove" that the duress defense is unpalatable as well? The concern is legitimate, but it turns out to be unfounded. Take a standard case of genuine duress: the defendant who perjures himself because he has been told he will be tortured if he does not. Let us do for this case what we did for the $10,000 "pseudo" duress case above: let us find a parallel case of temptation. Suppose the defendant is wracked by a terribly painful illness for which someone promises him a cure if he is willing to perjure himself. And suppose he does perjure himself. How do we feel about this parallel case? I think we would be strongly inclined to acquit, despite the fact that technically, we are dealing with a case of temptation and not duress. We would probably stretch the defense of duress so as to cover such a case, even if by its usual wording, it does not quite reach it. What this tells us is that there really is a deep difference between cases of valid and "pseudo" duress, a difference that is more than a matter of degree, a difference that warrants giving complete credit in the one kind of case and not even partial credit in the other kind. The difference is that if we try to "convert" a "pseudo" duress case into a case of temptation — as we did in the paragraph above — we end up with a situation in which we have no interest in showing the defendant any indulgence. By contrast, if we convert a genuine duress case into a case of temptation, we end up with a situation in which we are as inclined to acquit the defendant as we were when we framed it in terms of coercion.

Now this analysis of duress poses a bit of a problem for our understanding of the partial credit problem in negligence. Structurally, duress and negligence operate identically in denying partial credit. But the reasons for that in duress do not easily generalize to negligence. The duress argument depends critically on the presence of coercion and on our desire to treat certain cases of coercion and temptation symmetrically. Nothing like that seems to be going on with negligence. I do not see that as invalidating the parallel between the two domains, but it presents a puzzle.

Let us turn next to self-defense.

B. Self-Defense

Why is it that we do not grant a partial self-defense defense? We do something akin to that in jurisdictions that recognize imperfect self-defense:

we grant the defendant, who killed under a negligent misapprehension that he was acting in self-defense, a reduction by reclassifying his crime as manslaughter. But negligent misapprehensions are unusual in this regard.[12] Other kinds of imperfect self-defense are not treated in the same manner. Consider, for instance, the imminence requirement. A defendant is entitled to use self-defense only if he faces an imminent threat. That means, basically, that his attacker must have formed the intention to attack him and must have gone some ways toward implementing that intention. Imminence could easily be construed to be a matter of degree. There are different degrees of proximity that the attacker might have achieved to his ultimate goal. One might think that a defendant should get credit for his defensive actions in proportion to the proximity that the attacker has achieved to his attack. But that is not how it works. There is some point at which a valid self-defense argument is triggered, and prior to that, the defendant has no such right. On closer inspection, however, that abrupt transition makes a good deal of sense. It mirrors the law of attempt. We do not punish someone for an attempt until his intention is fully formed and until he has crossed the (admittedly hard-to-specify) threshold between preparation and attempt. Just as the criminal is not vulnerable to arrest until he has crossed that point, he is not vulnerable to defensive force until he does.

Take another aspect of self-defense. One may use force to retrieve stolen property so long as one does so in fresh pursuit of the thief. Thereafter one must depend on the authorities to do the retrieving. Suppose a defendant has exceeded that period. That is, he continues to hunt the thief even though he is by now in "stale" pursuit. One might be inclined to give him credit in proportion to the freshness or staleness of his cause, so that if he falls just short of meeting the fresh pursuit requirement, he will be found to be just a little more guilty than had he fallen within it. Why do we not proceed in this fashion? The fresh pursuit rule seems to be but a special case of a more general moral distinction, that between undoing and preventing. Perhaps the most intuitively compelling instance of that distinction is the Case of the Dithering Victim: Knife in hand, Attacker takes a deadly lunge at Victim. As it happens, Victim is fully prepared. He could easily save himself by shooting Attacker. But he chooses not to. He hates to kill unless he absolutely hates to, and he hopes that he will able to weather whatever injury Attacker ends up inflicting. Alas, Attacker's attack ends up inflicting

12 In addition, they contain within themselves the negligence puzzle we are trying to explain: Why is it that we do not give the negligent self-defender credit for the low-probability possibility that he was actually killing an attacker?

a mortal wound. Or rather, it will be a mortal wound unless Victim obtains a heart transplant. There being no other donor around, Victim proposes that he now be allowed to kill Attacker after all and confiscate Attacker's heart for his own use. "All I am doing," he reasons, "is causing Attacker's death so as to save my life. I delayed doing so as long as I could in the hope that I would be able to save my life without taking Attacker's, but that turns out not to be possible. Therefore I ask that you now allow me to kill him to enable me to survive." Obviously, he is not going to get his wish. The right of self-defense has expired. It expired at the point at which Victim's killing of Attacker would no longer constitute the *prevention* of the attack, but the *undoing* of its effects. The transition from prevention to undoing is a quick and radical one, and the right of self-defense evaporates in a correspondingly sudden and complete manner.

This analysis of self-defense both adds to and subtracts from the partial credit problem in negligence. It subtracts because it substantiates our feeling that there is nothing unique about our refusal to give such partial credit. Our refusal to recognize a partial version of self-defense amounts to the same thing. But it also adds to the mystery, because the reasons for not recognizing a partial version of self-defense seem so particular to self-defense. The arguments about attempts and about the prevention/undoing distinction do not really suggest analogies in the negligence area.

C. Necessity

Let us turn finally to the defense of necessity. Why is there no partial necessity defense, and can the reason for its absence be used to explain our approach to negligence? There is special reason here to be hopeful. Conceptually, necessity and negligence seem intimately related. Both seem to involve a weighing of certain costs and benefits. The Model Penal Code's definition of necessity provides that "conduct which the actor believes to be necessary to avoid a harm or evil to himself is justifiable [if] the harm or evil sought to be avoided by such conduct is greater than that sought to be prevented by the law defining the offense charged."[13] That sounds very similar to the Learned Hand approach to negligence. Indeed, the only material difference seems to be that necessity involves the infliction of harm that is certain, whereas negligence involves the infliction of harm that is uncertain.

So why is it that we decline to grant a partial necessity defense to a defendant seeking to avoid a harm or evil that is just a little less serious than

---

13 Model Penal Code § 3.02 (Proposed Official Draft 1962).

the harm sought to be prevented by the crime charged? In other words, why do we not calibrate his blameworthiness, and proportion his punishment, to the gap between the "harm or evil sought to be avoided" by the defendant and the "harm or evil sought to be prevented by the law defining the offense charged"? Let us see how one might go about doing that. The most obvious way of doing so would be to take the harm the defendant commits and just subtract it from the harm he is averting. In the simplest case, if he destroys $1000 worth of property to avert a $700 loss, his culpability would be proportional to $300 and identical to that of someone who destroys $2000 worth of property to avert a $1700 loss. So far that sounds pretty reasonable. It starts to run into a bit of a problem with examples like the following: the defendant commits perjury because thugs threaten to inflict property damage on him of about $700. How do we measure the harm of the perjury? Is it the wrongful verdict that the perjury precipitates? Is it the systemic cost to the adjudicative process inflicted thereby? Something else altogether?

There are more troublesome cases yet. Consider the kind of case in which the defendant kills someone to save several lives, typically illustrated by the surgeon who harvests the heart, lungs, and kidneys of a healthy patient to save five people in need of transplants. The harm committed in this case is the killing of one person. The harm averted is the death of five other people who would otherwise be permitted to die. If we simply subtract one from the other, that would suggest that the defendant is not guilty of anything, since he averted more deaths than he precipitated. But obviously that cannot be right. We do not grant the necessity defense to in such cases. The defense is *not* usually understood to underwrite simple utilitarian tradeoffs. But if we are not willing to simply net out harms in the transplant case, that probably means we are wrong to do it in other cases as well. For instance, imagine a surgeon who kills *five* healthy patients so as to use their organs to save *one* dying patient who needs multiple transplants. Should we give the surgeon partial credit for his saving of one life? Should we subtract that life he saved from the lives he took? Since we do not give *complete* credit when the surgeon kills one to save five, it seems odd to give him *partial* credit when he kills five to save one. Put differently, if we refuse to net things out in the first kind of case — if we refuse to say that the surgeon acted correctly because he achieved a net saving of four lives — then it seems we should also refuse to net things out in the second case; in other words, we should refuse to say that he is only guilty of four murders rather than five because he only inflicted a net loss of four lives.

Perhaps where we went wrong was in focusing so much on the *harms* the defendant commits and averts. In defining the necessity defense, the Model Penal Code somewhat ambiguously asks us to compare "the harm or evil

sought to be avoided" by the defendant with the "harm or evil" associated with the crime he commits. Perhaps we should pay more attention to the notion of evil than to that of harm. Perhaps we should not be comparing harms, but evils. The two are not the same. An attempted murder may produce no harm, but it does represent an evil. Perjury may not do much harm, but it represents an evil. One might therefore interpret the necessity defense to say that the defendant is innocent of a crime if the *evil* (rather than the harm) associated with the crime he committed is less than the evil he prevented. And if we wanted to create a partial necessity defense consistent with this new approach, we would then say that where the defendant commits a crime that is more evil than the thing he prevents, he should get partial credit for the evil he prevents. Thus the defendant who perjures himself to avert some property damage deserves credit for the property damage he averted and should only be punished for the difference between those two evils. Once we interpret the necessity defense in this way, it seems we are able to avoid those problematic transplant cases. We will be able to convict the transplant surgeon despite the fact that he effectuated a net saving of lives. The reason is that he produced a net increase in evil, because the evil of killing one outweighs the evil of letting five others die.

Unfortunately, upon reflection, this is not going to work. Consider a standard necessity case: The defendant comes across a stranger who has been badly injured and needs to get to a hospital. The defendant steals a nearby car to drive her there. Ordinarily, he would be guilty of a crime. Under the circumstances, necessity gets him off the hook. But now consider how this case would be analyzed under an evil-based approach. We would be asking whether the evil the defendant committed — the theft — is more serious than the evil he prevents: letting the stranger die. Theft is a crime. Letting a stranger die is not. Thus engaging in the former is more of an evil than engaging in the latter. We could try to get out of this difficulty by saying that letting the stranger die is more evil than stealing the car. But in that case, why do we punish people for stealing cars but not for letting others die?[14]

Is there some other way of reinterpreting what we are comparing when we apply the necessity defense, so that we get a more acceptable set of answers? That seems unlikely for the following reason: we know that there are many different kinds of cases involving the taking of some lives to save a larger number. Some of those cases — like the transplant scenario — are ineligible

---

14 Frances M. Kamm draws attention to this profound paradox in her path-breaking article *Supererogation and Obligation*, 82 J. Phil. 118 (1985).

for the necessity defense, but certain others are eligible for it. For instance, the well-known trolley case. What distinguishes the trolley case from the transplant case has solely to do with the manner in which the one is sacrificed for the sake of the five. What it is about the manner in which this occurs in each case that makes them morally different continues to be a matter of some dispute. Judith Jarvis Thomson famously proposed her "redistribution principle" to account for the difference.[15] Michael Moore has suggested an explanation having to do with degrees of causation: the transplant surgeon constitutes a "big" cause of the victim's death, whereas the trolley driver only a "small" cause.[16] Frances Kamm has offered a proposal that I find the most persuasive to date. She thinks that some variant of the double effects doctrine is at work, something she calls the principle of permissible harm, which says that one can harm a smaller number of save a larger number just as long as the "saving" is more proximate than the "harming."[17] In other words, one may kill several if it is the further outgrowth of presently saving many more; but one is not allowed to do the converse. To put it in a third way yet, it is all right to act so as to primarily save the several and secondarily kill the one, but one may not act so as to secondarily save the several and primarily kill the one. The trolley driver is primarily saving the five and secondarily killing the one. The transplant surgeon is secondarily saving the five and primarily killing the one.

If I am right in concluding that the necessity defense is not a matter of netting out, then it becomes clear why it would be hard to create a partial version of it. If necessity simply were to involve netting out two things and seeing whether one thing (e.g., the harm prevented) is bigger than another (e.g., the harm done), then we could say that a partial necessity defense exists when the one thing (the harm prevented) is somewhat smaller than the other (the harm done). But if necessity were not a matter of netting out, what would a partial version of it look like? What, for instance, would a partial version of the distinction between the transplant case and the trolley case look like? Should we say that a case that somehow lies halfway between

---

15 A trolley is heading down an incline when Edward, its driver, discovers that the brakes are not working. On the track ahead of him are five people: the banks are so steep that they will not be able to get off the track in time. The track has a spur leading off to the right, and Edward can turn the trolley onto it. Unfortunately, there is one person on the right hand track. Edward can turn the trolley, killing the one; or he can refrain from turning the trolley, which would mean the death of the five. Judith Jarvis Thomson, *The Trolley Problem*, 94 Yale L.J. 1395, 1935 (1985).

16 Michael S. Moore, Placing Blame 701 (1997).

17 Frances M. Kamm, Morality, Mortality II pt. 2 (1996).

a trolley scenario and a transplant scenario should be treated as a case of partial necessity? That too is not likely to work.

Nevertheless, many people will be left with an abiding sense that there must be something to the "partial necessity" idea. If we compare two robbers, one of whom acted out of greed and the other of whom was trying to raise funds for his ailing mother's cancer treatment, is it not clear that the second is better than the first? To be sure, we would not want to acquit either of them, but if we do think the second is better, does that not indicate that we view him as partially justified? No, I do not think it does. To be sure, there is an interesting issue concerning the extent to which we should distinguish between worthier and unworthier motives that might be driving the defendant. But even if we were to think that worthier motives make a defendant less culpable, that does not mean we would think him partially justified. It would be strange to describe a motive that is less unworthy than some other motive as being partially justifiying. To pick an extreme case, surely it is less unworthy to kill out of greed than out of pure sadistic pleasure; but just as surely, this is not because greed is more of a justification than sadistic pleasure. Moreover, it is not altogether clear that the worthiness of motives ought to have any bearing at all on someone's blameworthiness. There are many aspects of someone's misconduct that speak to his general moral worthiness, without having any bearing on his criminally relevant blameworthiness. Compare the proverbial Bad Samaritan — the malicious bystander who could easily pull the drowning baby out of the pond, but refuses — with a pickpocket. We have no trouble blaming and punishing the pickpocket. We are disinclined to punish the Bad Samaritan. Presumably because we do not consider him blameworthy — at least not in the sense relevant to the criminal law. Of course, there is *some* sense in which he is blameworthy, which is why we are willing to condemn him, probably even more harshly than the pickpocket. But that does not seem to be the sense of blameworthiness that is relevant to the criminal law. In other words, the Bad Samaritan's refusal to carry out an easy rescue obviously reflects on his general moral worthiness, but not, it seems, on his criminally relevant blameworthiness. The robber's motives may well be of the same character: they may reflect on his general moral worthiness, but not on his criminally relevant blameworthiness.

Let us return then to the question about negligence that prompted this whole inquiry. Do our reasons for not giving partial credit for partial necessity explain why we do not give partial credit for negligence? Superficially, necessity and negligence seem very much alike. Both necessity and negligence are commonly described as involving a netting-out of certain things. In the case of necessity, we are supposedly netting out the harm

done and the harm prevented; in the case of negligence, the risk incurred and its justification. If the risk incurred is 100% certain (if it becomes a harm done rather than a harm risked), the two concepts seem to actually collapse into each other. On closer scrutiny, however, it starts to look as though the two concepts are in fact quite different. As we just saw, we generally find it intolerable to decide whether someone merits the necessity defense by simply netting out the harm he did and the harm he avoided. The same does not seem to be true when we decide whether someone has behaved negligently. Here we often find it perfectly tolerable to net out the probabilistic harm the defendant risked and the benefit he thereby achieved and to decide on the basis of that net result whether care was adequate or not. It seems we are willing to be far more consequentialist in judging the infliction of uncertain harm as opposed to certain harm. To make this more concrete, think of a driver who is racing to deliver a group of critically ill patients to the ER. Suppose he sees a single person on the road ahead of him and rightly calculates that if he swerves or slows down to avoid hitting him, enough time will be lost to result in the death of everyone in his car. May he speed ahead and run over the person on the road? No, of course not. This is the classic case of killing one to save several. Just the sort of consequentialist calculation we are not permitted to indulge in. Now let us change the facts a little. Suppose that in calculating his prospects, he is somewhat more uncertain than before. He does not flatly predict that if he swerves or slows down, this will necessarily entail the death of his passengers. Rather, he believes that there is a 50% chance that it will result in their deaths. Nor, for that matter, does he flatly predict that if he does not swerve, that will necessarily mean running into the person in front of him. Rather, he believes that there is a 50% chance that it will mean that. In other words, we have reduced both the probability of the harm and the probability of the benefit that harm might generate by 50%. If he decides to keep going, to neither swerve nor slow down, we might well judge him not to be acting negligently — since his expected benefit from keeping going exceeds his expected cost from doing so. This is what I mean when I say that consequentialist calculations that are unacceptable in the assessment of necessity are perfectly acceptable in the assessment of negligence.

But we are falling prey here to an illusion. In fact, we are no more consequentialist in our assessment of risk-taking than in our assessment of necessity. There is an important difference between the two versions of the ER-hypo I posed above, an important difference *in addition* to the difference between certain and uncertain harms. *The driver in the certain case is primarily killing the person he is running over and only secondarily saving the people he is driving to the ER.* His saving of the latter is a further

outgrowth of running someone over. *By contrast, the driver in the uncertain case is only secondarily killing the person he is running over and primarily saving the people he is driving to the ER.* What he is most immediately accomplishing by speeding is the saving of his passengers; the running over of the one is an unfortunate, further outgrowth of that activity. The driver in the uncertain case is, therefore, quite unlike the garden-variety necessity case, and it should not astonish us that he invites different treatment. But although the driver in the uncertain case does not resemble the garden-variety necessity case, he does resemble a certain more unusual kind of necessity case, namely, the trolley scenario. The trolley driver too is someone who saves primarily and kills secondarily. When evaluating the trolley driver, we are happy to be consequentialist. We grant him his defense if more are saved than die. And that is exactly what we do when we assess negligence.

So what can be learned from our analysis of necessity about the question of whether to credit the negligent defendant with the benefit he conferred? We now know that negligence closely tracks one particular strain of necessity, the trolley scenario. What is more, we are fairly confident that in the trolley scenario, we would not give someone partial credit who turned the trolley so as to kill the five rather than the one. We would not find that person to be guilty merely of four murders rather than five, on the ground that the saving of one cancels out one of the five killings. That does not directly answer the partial credit problem about negligence, but, rather, it turns it into a different problem: Why is it that we do not give the trolley driver who kills five credit for the one he saved?

Where does this leave us? Our treatment of defenses remarkably mirrors our treatment of negligence. In both cases, we engage in a comparison of sorts. Negligence requires us to compare the harm risked with the justification for risking it. The defenses require us to compare the crime committed with the justification or excuse for committing it. Once we have determined that there is negligence, we then forget about the justification and focus solely on the harm risked in assessing blameworthiness. Once we have determined that there is no valid defense, we then forget about the defense and focus solely on the seriousness of the crime committed. It is hard to believe that there are not analogous subterranean reasons at work. And yet we had trouble finding them. The only close analogy at all was between negligence and a certain type of necessity case, the trolley scenario. And that analogy was only partly helpful because we are not sure why it is that we do not have a partial necessity defense. Nevertheless, we are, I think, entitled to take presumptive comfort from the similarity between negligence and the defenses. The fact that they exhibit identical brands of weirdness,

and that we are able to give good reasons for the weirdness with regard to defenses, does, I think, significantly support our treatment of negligence.

V. PARTIAL CREDIT AND THE OFFENSES

Although the workings of the defenses most clearly mirror the "no partial credit" feature of negligence, another set of analogies is arguably more in point. Negligent harming is just one kind of criminal offense. It ought therefore to be instructive as to whether other offenses raise the same partial credit problem and how they handle it.

A. Fraud

Let us begin with crimes of deception. The quintessential fraud involves a defendant who tells his victim a lie, which then leads the victim to part with some valuables. How badly has this defendant behaved? That will depend on the value of the valuables. What it will not depend on is the mendacity of the lie. To be sure, the lie might be too slight to trigger liability. Mere puffery will not suffice. But suppose the lie is more serious than that, serious enough to constitute a criminal offense. Once that threshold has been crossed, we do not care how lying the lie was, how far-off from reality the defendant's misrepresentations swerved. In this regard, fraud is very much like negligence. The mendacity of the lie counts in determining whether there was fraud, but not in determining how serious that fraud was. We do not give partial credit for "comparative lack of mendacity." However much short the liar's lie stops of being maximally egregious, we do not really care, just as long as it was egregious enough to count as fraudulent. In that regard, then, fraud has the same structure as negligence.

B. Coercion

Consider next crimes of coercion. Again, take a look at some quintessential cases. The defendant tells her victim that unless he pays her $10,000, she will:

Case 1: kill him;
Case 2: punch him;
Case 3: strip him of all his clothes;
Case 4: spread false rumors about him;
Case 5: spread true, but embarrassing rumors about him;
Case 6: cease to be his friend.

Suppose now that in each of these cases, the victim pays the defendant the requested sum. As we move from Case 1 to Case 6, the defendant's threatened misconduct vis-à-vis the victim progressively decreases in moral seriousness: killing is more serious than punching, which is more serious than stripping the victim of all his clothes, which is more serious than spreading false rumors, which is more serious than spreading true but embarrassing rumors, which is more serious than ceasing to be his friend. Our intuition about the relative moral seriousness of these various kinds of misconduct is supported by the law's treatment of them. Killing is a more serious crime than punching, which is a more serious crime than stripping someone of all his clothes. Spreading false rumors is no crime at all, but a serious tort. Spreading true rumors is a tort in some jurisdictions (depending on further circumstances such as whether it constitutes an invasion of privacy), but not in others, which squares with our sense that it is less serious than spreading rumors that are false. Ceasing to be someone's friend is neither a crime nor a tort, and you may be wondering why I consider it misconduct at all. The reason I do is that I am assuming that the cessation of the friendship is in retaliation for not paying $10,000. That does not make it a legally sanctionable (i.e., punishable or compensable) piece of misconduct, but misconduct nevertheless.

Notice, however, the following intriguing oddity. Although the threatened misconduct decreases as we move from Case 1 to Case 6, the overall blameworthiness of the defendant's behavior does *not* decrease commensurately. Let me clarify what I have in mind. Compare Case 1 and Case 2. In other words, compare "Pay me $10,000 or I will kill you!" with "Pay me $10,000 or I will punch you!" followed in each case by the surrender of $10,000. Both of these are straightforward cases of robbery, of roughly equal seriousness. They are equal despite the fact that the underlying threats — to kill versus to assault — are quite unequal. Both cases involve the forced extraction of the victim's property, and that makes them equal, notwithstanding the unequal seriousness of the threatened behavior used to extract that property. Compare next Case 2 and Case 3, i.e., compare "Pay me $10,000, or I will punch you!" with "Pay me $10,000, or I will spread false rumors about you!" Case 2, we already noted, is a straightforward case of robbery. It is not entirely clear whether Case 3 would be classified as robbery or blackmail, but either way, it would be considered of roughly equal seriousness as Case 2. Case 2 and Case 3 are equal also despite the fact that the underlying threats — to assault versus to spread false rumors — are quite unequal. They are equal because both involve the forced extraction of property; the inequality of the threats is irrelevant. By the same token, Case 4 is of roughly equal moral seriousness as Case 3 and Case 5 of roughly

equal seriousness as Case 4. Case 6 is a different story. Case 6 is no crime at all.

Step back now and take a look at the total picture that emerges. As we gradually diminish the underlying threats involved in Cases 1 through 6, the overall blameworthiness of the actor does not diminish until we move from Case 5 to Case 6. As between those last two cases, a relatively slight decrease in the seriousness of the threatened behavior translates into a major change in overall blameworthiness. Up until that point, even very significant decreases in the seriousness of the threatened behavior — such as the decrease between Case 1 and Case 5 — hardly translates into any decrease in overall blameworthiness. In other words, despite the fact that the defendant reduces the moral seriousness of the misconduct she threatens as we move from Case 1 to Case 5, she receives no partial credit for that. Our assessment of the overall seriousness of her behavior stays roughly constant, until, that is, the threatened misconduct has become sufficiently inoffensive — as it does in Case 6 — that we decide not punish it at all. This, of course, rather exactly mirrors negligence. Up until the risk imposition reaches the level of unjustifiability, there is no criminal liability. And past that threshold, the liability does not vary much with the level of justifiability, but only with the level of risk imposed.

C. Treason

The quintessential case here is the trial of William Joyce, a.k.a. Lord Haha, for broadcasting anti-British propaganda, addressed to British servicemen, over the German radio waves during the Second World War. The broadcasts became quite infamous for their sneering, relentlessly derisive calls to British soldiers to lay down their arms and recognize the insurmountable superiority of German might, virtue, and culture.[18]

When he was tried at the Old Bailey after the War, Joyce argued that he could not be guilty of treason because he was neither a British citizen nor someone who owed allegiance to Britain on some other well-established ground, such as being a non-citizen permanent resident. The issue of Joyce's citizenship and allegiance was, indeed, a murky one, because Joyce had been born in the United States to a naturalized Irish immigrant who subsequently took his family back to England, but without surrendering their American citizenship. On the face of it, that would have seemed to make it impossible to prosecute Joyce.

---

18 Joyce v. D.P.P., 1946 A.C. 347.

Recognizing this difficulty, the prosecution tried to nail Joyce on the somewhat more obscure ground that he owed the kind of allegiance to the British Crown that is imposed even on non-citizens while they reside in Britain. Joyce, of course, was no longer residing in Britain when he made his broadcasts, but some hoary precedents had held that under certain circumstances, the allegiance owed might continue

> if such alien [while continuing to maintain] a Family and Effects here should during a War with his Native Country go thither and there Adhere to the King's Enemies for the purpose of Hostility [.] He might then be dealt with as a Traitor. For he came here under the Protection of the Crown. And though his Person was removed for a time, his Effects and Family continued still under the same Protection.[19]

Joyce, however, had moved not only himself, but his entire family and all of his possessions to Germany. That left the prosecution with only one very tenuous ground for "allegiance": Joyce had taken out a British passport, to which he was not in fact entitled, since he was not a British citizen, but the authorities did not at the time realize that. He exited Britain using that fraudulently obtained passport. The court agreed that traveling on this passport was enough, although just barely enough, to trigger the allegiance that was the necessary foundation of a treason charge. On this slender reed rested Joyce's conviction and his subsequent hanging. Rebecca West thought that he

> die[d] the most completely unnecessary death that any criminal has ever died on the gallows. He was the victim of his own and his father's lifelong determination to lie about their nationality. For had he not [obtained that] English passport and had he left England for Germany on the American passport which was rightfully his, no power on earth could have touched him. As he became a German citizen by naturalization before America came into the war, he could never have been the subject of prosecution under the American laws of treason.[20]

Considered in this light, Joyce's actions appear a very minor sort of treason. Treason that just barely qualifies as treason, because he traveled on a British passport. But if we judge the matter thus, it is because we are giving the defendant "partial credit" for almost but not quite managing to avoid committing treason. The law does not in fact proceed in this way. It

---

19 Rebecca West, The Meaning of Treason (1947).

20 *Id.*

gives no such partial credit. Once it has determined that allegiance is owed, however partially justifiable it seems for him not to render it, it measures the wickedness of the treason solely by the extensiveness of the defendant's participation in the enemy country's war effort. In short, the law proceeds here too as it does in negligence.

D. Murder

On certain occasions, the defendant's liability for murder familiarly depends on whether what he did is best characterized as a killing or a letting-die. Absent special duties, only the former, but not the latter, constitutes murder. To determine which characterization is the right one will often require us to locate certain controversial baselines. Think of the defendant who has thrown a rope to some drowning swimmers, but yanks it back as the swimmers are about to reach for it. Whether this is a case of killing or letting die has to do with whether we think that the rope floating in the swimmers' direction constitutes a new status quo, a new baseline, with which the defendant is interfering or whether we think that the absence of the rope continues to constitute the status quo, the morally relevant baseline. If it is the former, the defendant is guilty of murder; if it is the latter, he is not. Choosing between those two alternatives requires us to weigh numerous pros and cons, competing analogies, symmetries, principles, and policies. The gap between these pros and cons in a given case might be a narrow or a wide one. The more controversial the case, the narrower the gap of course. But the defendant's blameworthiness for murder is not proportional to that gap. The defendant's blameworthiness will be a function of the number of people killed and of various other "general part" factors (such as whether he acted with intention or with mere recklessness, whether he was a principal or an accomplice, and so forth). What it will not be a function of is *how close* he came to having committed a mere letting-die rather than a killing. We do not give partial credit for having made it part of the way toward "letting-die" status. The parallel to negligence is evident.

VI. Lessons

Our refusal to give the negligent actor partial credit thus seems to track what we do in the criminal law generally. We do not ask how fraudulent the defrauder's lie was, but merely whether it was fraudulent enough, and if it was, we measure blameworthiness by focusing on the harm done. We do not ask how coercive the robber's threat was, but merely whether it was

coercive enough, and if it was, we measure blameworthiness by focusing on the harm done. We do not ask how strong the traitor's ties to his country are or how strong the case for his owing allegiance, but merely whether they were strong enough, and if they were, we measure blameworthiness by the extent of his participation in the enemy's war effort. We do not ask how strong the case for drawing the baseline by which we evaluate the murderer's actions, but merely whether it was strong enough to draw it where we did, and then measure his blameworthiness by focusing on the harm done. The evaluation of risk-creating conduct seems entirely of a piece with all of this. Our refusal to give partial credit simply means that we do not ask how unjustifiable the defendant's risk creating conduct was, but merely whether it was unjustifiable enough, and if it was, we measure blameworthiness by focusing in the harm risked, or perhaps the harm done, but not its partial justifiability.

But why do we do things the way we do with regard to these other offenses? Are there good reasons, and if there are, do they really extend to the evaluation of negligence? When exploring the analogy between negligence and the defenses, it seemed as though many of the reasons for not recognizing partial defenses are, in fact, quite peculiar to those defenses; perhaps the same is true here as well? I suspect not. The structural similarities just seem too strong not to suppose that a common logic is at work. What that common logic is will bear further looking into.[21]

---

21 I doubt that it has much to do with practical considerations, like the need to draw clear lines to save administrative costs. Our assessment of blameworthiness simply reflects moral judgments. Practical considerations of this sort may well play a role in deciding how prominent a role to accord blameworthiness in meting out punishment. But that is a separate issue.