

[This essay originally appeared in *Ethics in Practice*, edited by Hugh LaFollette (Oxford: Basil Blackwell, 1997), pp. 470-479.]

Punishment and Desert

James Rachels

When someone who delights in annoying and vexing peace-loving folk receives at last a right good beating, it is certainly an ill, but everyone approves of it and considers it as good in itself, even if nothing further results from it.

—Immanuel Kant, *Critique of Practical Reason* (1788)¹

Retributivism—the idea that wrongdoers should be “paid back” for their wicked deeds—fits naturally with many people’s feelings. They find it deeply satisfying when murderers and rapists “get what they have coming,” and they are infuriated when villains “get away with it.” But others dismiss these feelings as primitive and unenlightened. Sometimes the complaint takes a religious form. The desire for revenge, it is said, should be resisted by those who believe in Christian charity. After all, Jesus himself rejected the rule of “an eye for an eye,”² and St. Paul underscored the point, saying that we should not “return evil for evil” but we should “overcome evil with good.”³ To those who adopt this way of thinking, whether on secular or religious grounds, vengeance cannot be an acceptable motive for action.

This objection is, for the most part, misguided. The idea that wrongdoers should be “paid back” for their wickedness is not merely a demand for primitive vengeance. It is

¹ Immanuel Kant, *Critique of Practical Reason*, translated by Lewis White Beck (Chicago: University of Chicago Press, 1949; but originally published in 1788), 170.

² Matthew 5: 38-41.

³ Romans 12: 17, 21.

part of a moral view with a subtle and complicated structure, that can be supported by a surprisingly strong array of arguments. The key idea is that people deserve to be treated in the same way that they choose to treat others—thus, those who treat other people badly deserve to be treated badly in return. Retributivism is just the application of this idea to the special case of punishment. If this were only a view about punishment, it would not be very compelling. But the idea that people should be treated according to their deserts is a central component of our general moral understanding, with applications in many different areas of life. If we were to stop thinking of people as deserving or undeserving of special treatments, our moral outlook would be unrecognizably different.

In what follows we will begin by looking at the general idea of desert. First, we will consider what it is like to treat people as they deserve. We can describe the practice of “treating people as they deserve” without making any judgments about whether it is a good thing. Thus we will begin by asking:

1. What does the practice of “treating people as they deserve” involve? What, if anything, makes someone deserving of special treatment?

Then we will turn to the normative question:

2. Should we treat people as they deserve? Are there any compelling reasons in favor of such a practice?

Finally, with this much in hand, we will turn to the special case of punishment.

1. Treating people as they deserve

Desert and past actions. Moral judgments about what to do frequently depend on considerations about what will happen as a result of our actions. These are future-oriented considerations. Thus, we should support the United Nations Children’s Fund because it will save lives, and we should stop eating meat because it will spare animals pain. Sometimes, however, moral judgments focus on the past: if you say that you ought

to do such-and-such because you promised, you are citing a fact about the past, not the future, as the basis of your judgment.

Desert is a backward-looking consideration. What a person deserves always depends on what he or she has done in the past. (The familiar lament, “What have I done to deserve this?” is not just an idle remark. When desert is at issue, it asks exactly the right question.) Consider this case:

The two candidates for promotion. The owner of a small business must decide which of two employees to promote. The first is a man who has been a loyal and hard-working member of the staff for many years. He has frequently taken on extra work without complaint, and in the company’s early days, when its future was in doubt, he would put in overtime without demanding extra pay. His efforts are one reason that the company survived and prospered. The other candidate is a man who has always done the least he could get by with, avoiding extra duties and quitting early whenever he could. To have convenient names, we may call them Worker and Slacker. Whom should the owner promote?

Clearly, Worker deserves the promotion. He has worked for it, and earned it, in a way that Slacker has not.

Deserving the promotion is not the same as needing or wanting it. Both Worker and Slacker might benefit from the promotion; perhaps both could use the extra money and the status it would bring. But this has nothing to do with desert. Although Slacker might benefit just as much from being promoted, he does not have the same claim to it as Worker because he has not earned it in the same way.

Nor is the question of desert the same as the question of who would perform better in the new position. Obviously there is reason to think Worker would do better, because he has shown himself to be more diligent in the past. But again, that is not the basis of Worker’s claim. Even if we knew that Slacker would reform and do just as well in the new position—the promotion may be just the prod he needs—Worker would still have an independently valid claim on the promotion, based simply on his past performance.

Does anything other than past performance affect what a person deserves? Sometimes it is assumed that people deserve things because of their superior native endowments. Thus, if Slacker is naturally smarter or more talented than Worker, it might be suggested that this makes him deserving of the promotion. This sort of idea was once commonplace, but it is no longer very popular among those who systematically study ethics. It has fallen into disrepute since the publication of John Rawls's *A Theory of Justice*. Rawls writes that:

Perhaps some will think that the person with greater natural endowments deserves those assets and the superior character that made their development possible. Because he is more worthy in this sense, he deserves the greater advantages that he could achieve with them. This view, however, is surely incorrect. It seems to be one of the fixed points of our considered judgments that no one deserves his place in the distribution of native endowments, any more than one deserves one's initial starting place in society.⁴

Rawls refers to "our considered judgments," but there is something more here than an appeal to our beliefs. There is also an implicit argument, namely that native endowments are not deserved because no one *does anything* to deserve them: they are a result of a "natural lottery" over which we have no control. If you are naturally smarter or more talented than other people, then you are just luckier; and you do not deserve better merely on that account. This fits well with the idea that people deserve things because and only because of their past actions.

What else might plausibly be thought to provide a basis for desert? It has sometimes been suggested that *achievements* are pertinent. It may be argued that Slacker could deserve the promotion, despite Worker's greater effort, if Slacker had succeeded in contributing more to the company. (If he were smarter and more talented than Worker, for example, Slacker's efforts might have produced more despite his having exerted himself less.) On reflection, however, it does not appear that the reference to

⁴John Rawls, *A Theory of Justice* (Cambridge, Mass.: Harvard University Press, 1971), pp. 103-4.

achievements adds anything substantial to the discussion. Achievements are the products of native endowments combined with effort—often with a good bit of luck thrown in—and if one cannot deserve things because of one’s native endowments, neither can one deserve things because of the achievements that those endowments make possible. To see what someone deserves we have to separate the two components (native endowments and work) and identify the contribution made by each. The maximally deserving man or woman is not simply the one who achieves the most, but the one who achieves the most he or she can given the abilities with which he or she is endowed. The key idea, as far as desert is concerned, is “doing the best you can with what you have.”

Moral deserts. Moral deserts are deserts that one has, not in virtue of one’s performance in a special type of activity (such as working at a job), but in virtue of one’s more general way of dealing with other people. We can choose to treat others well or badly, and what we deserve from them in return depends on the choice we make.

Consider this example:

The ride to work. Suppose you, Smith, and Jones all work at the same place. One morning your car won’t start and you need a ride to work, so you call Smith and ask him to come by for you. But Smith refuses. It is clear that he could do it, but he just doesn’t want to be bothered, so he makes up some excuse. Then you call Jones, and he gives you the ride you need. A few weeks later, as you are about to leave for work, you get a call from Smith. Now he is having car trouble and he asks you for a ride. Should you accommodate him?

Perhaps you will think you should help Smith, despite his own unhelpfulness—after all, it would be little trouble for you, and by helping him you might teach him a lesson in generosity. But if we focus on what he *deserves* a different answer seems obvious: he deserves to be left to fend for himself. Jones, on the other hand, is an entirely different matter. If Jones should ask you for a ride, you have every reason to give it: not only will it help him, he deserves it. This is especially clear when we consider the case of a forced choice:

The simultaneous requests. Smith calls and asks for a ride. Meanwhile, Jones is on the other line also needing a ride. But they live in opposite directions, so that it is impossible for you to help both. Which do you help?

If we did not concede that Jones's past conduct makes him more deserving, we would be hard put to explain why it seems so obvious that helping Jones is the mandatory choice.

Particular people may be especially obligated in this way. If someone has done you a favor, *you* are indebted to them and you specifically owe it to them to return the favor. It is you, and not someone else, who owes Jones the ride. Sometimes this is thought to end the matter: if someone has helped you, it is said, you are indebted to them; otherwise you have no obligation. But it is short-sighted to view things in this way. Anyone at all can justifiably take it as a good reason for treating someone well if that person has treated others (not necessarily oneself) well. Suppose Jones is habitually helpful to people, while Smith is not; but you personally have never had much interaction with either of them. Now suppose you must choose which to help, and you cannot help both. Surely their respective histories is a reason, even for you, to prefer Jones.

Thus we have:

The Principle of Desert: People deserve to be treated in the same way that they have (voluntarily) treated others. Those who have treated others well deserve to be treated well in return, while those who have treated others badly deserve to be treated badly in return.

The qualification "voluntarily" is included because if an action is coerced, accidental, done from ignorance, or the like, then the agent may not deserve anything on account of it.

The Principle of Desert has both a positive and a negative side. Those with a generous temperament may find one appealing but recoil from the other: they may like the idea that some people deserve good treatment but dislike the companion idea that others deserve bad treatment. After all, it seems ungrateful to say that someone who has treated others well does not deserve to be treated well in return; but to say that someone deserves to be treated badly seems, on the face of it, mean-spirited and unsympathetic.

So it might be suggested that we keep the idea of positive desert but discard the idea of negative desert.

But this won't work. If we jettison one we will surely have to jettison the other along with it. Superficially it may appear that we could split them apart: we could *say* that some people deserve good treatment but that no one ever deserves ill, and if we go no further this might seem consistent. But the inconsistency would emerge when we tried to provide a rationale for this combination of beliefs. What reasoning could justify holding that good performance merits a positive response that would not also imply that bad performance merits a negative response? The answer, so far as I can tell, is none.

2. Why people should be treated as they deserve

So far we have merely described, in a rough-and-ready way, what it means to treat people as they deserve. But we have given no reason whatever for thinking this is a good thing. Should we treat people as they deserve? Or, having seen what the practice involves, should we reject the whole business? There are at least three reasons for treating people according to their deserts. These add up to a compelling case.

First, acknowledging deserts is a way of granting people the power to determine their own fortunes. Because we live together in mutually cooperative societies, how each of us fares depends not only on what we do but on what others do as well. If we are to flourish, we need to obtain their good treatment. A system of understandings in which desert is acknowledged gives us a way of doing that. Thus, if you want to be promoted, you may earn it by working hard at your job; and if you want others to treat you decently, you can treat them decently.

Absent this, what are we to do? We might imagine a system in which the only way for a person to ensure good treatment by others is somehow to coerce that treatment from them—Worker might try threatening his employer. Or we might imagine that good treatment always comes as charity—Worker might simply hope the employer will be nice to him. But the practice of acknowledging deserts is different. The practice of acknowledging deserts gives people control over whether others will treat them well or badly, by saying to them: if you behave well, you will be *entitled* to good treatment from

others because you will have earned it. Without this control people would be in an important sense impotent, less able to affect how others will treat them and dependent on coercion or charity for any good treatment they might receive.

I believe this is the deepest reason why desert is important, but there are others. A second reason is connected with the egalitarian idea that social burdens and benefits should be equally distributed. In working harder, Worker had to forgo benefits that Slacker was able to enjoy. While Worker was spending his time working, Slacker was able to do things that Worker might have liked to do but was unable to do. (This, of course, will be typical of any situation in which one person chooses to expend time and effort on a disagreeable task, while another person—faced with the same choice—opts for a more enjoyable alternative.) This suggests a simple argument for rewarding the harder worker: other things being equal, burdens and benefits should be distributed equally. Slacker has had a benefit (more leisure time) that Worker has not had, while Worker has had a burden (more work) that Slacker has not had. Giving Worker a benefit now (the promotion) may therefore be seen as nothing more than righting the balance. Contrary to superficial appearances, then, giving Worker the promotion does not make their respective situations less equal. On the contrary, it alters the situation in the direction of greater equality. This is a reason why even egalitarians might favor treating people according to their deserts.

These arguments apply equally well to moral deserts. Acknowledging moral desert permits people, who are after all largely dependent for their welfare on what other people do, to control their own fates by allowing them earn good treatment at the hands of others. They do not have to rely on coercion or charity. Moreover, those who treat others well will have, in the course of doing so, forgone benefits for themselves. There are costs involved in helping others—in giving you the ride Jones was inconvenienced in a way that Smith was not. So once again, reciprocating is a way of making the distribution of burdens and benefits more nearly equal.

To these reasons a third may be added. Morality includes (some would say it consists in) how we choose to treat other people in our myriad interactions with them. But if reciprocity could not be expected, the morality of treating others well would come to occupy a less important place in people's lives. In a system that respects deserts, someone who treats others well may expect to be treated well in return, while someone

who treats others badly cannot. If this aspect of moral life were eliminated, morality would have no reward, and immorality would have no bad consequences; so there would be less reason for one to be concerned with it. If people were perfectly benevolent, of course, such incentives would not be needed. But for imperfectly benevolent beings such as ourselves the acknowledgment of deserts provides the reason for being moral that is required for the whole system to be effective.

Punishment

Retributivism is the application of the Principle of Desert to the special case of criminal punishment: it is the view that people who commit such crimes as murder and rape deserve to be punished and that this alone is sufficient to justify punishing them. It is not merely that punishing them satisfies certain sorts of vengeful feelings. On the contrary, it is a violation of justice if murderers and rapists are allowed to walk away as if they had done nothing wrong. It is a matter of justice for the same reason that promoting Worker is a matter of justice or that preferring to help Jones rather than Smith is a matter of justice.

As we have seen, acknowledging deserts is part of a moral system that allows people, by their own behavior, to determine how others will respond to them. Those who treat others well elicit good treatment in return, while those who treat others badly provoke ill treatment in return. That is why, when a criminal is punished, it may be said that “He brought it on himself.” The argument concerning equality is also commonly invoked when punishment is at issue. There are costs associated with law-abidance. Law-abiding people have accepted a burden—inconvenient constraints on their conduct—which the lawbreaker has not accepted. Meanwhile the lawbreaker has had benefits denied to others (presumably, his illegal conduct was undertaken because there was profit of some kind in it for him). Punishment corrects things in the direction of greater equality. That is why it is commonly said that crime “upsets the scales of justice” and that punishing wrongdoers “restores the balance.”

But the charge that retributivism is a mere rationalization of vengeful feelings is not the only objection that has been made against it. Philosophers have faulted

retributivism on other, weightier grounds. Bentham, who believed that social policies should promote the general welfare, noted that retributivism approves of *increasing the amount of suffering in the world*—if a miscreant harms someone, and we “pay him back” by harming him in return, we have only added to the total misery. Bentham did not see how this could be right. In his *Principles of Morals and Legislation*, published only one year after Kant’s remark about the “right good beating,” Bentham wrote: “All punishment is mischief: all punishment in itself is evil.” Therefore, he concluded, “if it ought at all to be admitted, it ought only to be admitted in as far as it promises to exclude some greater evil.”⁵

How can punishment “exclude some greater evil”? The obvious answer is by preventing crime. In general, sanctions are needed to secure compliance with social rules. If there were no rules against murder, assault, and theft, no one would be safe; we would live in a Hobbesian State of Nature in which life would be “solitary, poor, nasty, brutish, and short.” To avoid this, it is not enough to ask people politely if they would mind behaving themselves. Murder, theft, and the like cannot be left as matters of individual discretion. So, to ensure compliance with such rules, we attach sanctions to them. We do not say “Please do not murder”; we say “You *must not* murder, *or else*.”

The idea that punishment is justified as a means of preventing crime is so natural and appealing that we might expect it would dominate social-scientific thinking about the criminal justice system. Surprisingly, however, it has not been particularly influential. During the past 150 years a different sort of conception has prevailed. In the latter half of the nineteenth century it was argued that, if we are serious about preventing crime, we should try to identify its causes and do something about them. Crime, it was said, results from poverty, ignorance, and unemployment; therefore, social energy should be directed toward eliminating those blights. Moreover, when individuals commit crimes, rather than simply punishing them, we should address the problems that caused their aberrant behavior. People turn to crime because they are uneducated, lack job skills, and have emotional problems. So they should be educated, trained, and treated, with an eye to making them into “productive members of society” who will not repeat their offenses. In

⁵ Jeremy Bentham, *An Introduction to the Principles of Morals and Legislation* (New York: Hafner Publishing Company, 1948; but first published in 1789), 170.

enlightened circles this came to be regarded as the only sensible approach. As Bertrand Russell once put it,

No man treats a motorcar as foolishly as he treats another human being. When the car will not go, he does not attribute its annoying behavior to sin; he does not say, “You are a wicked motorcar, and I shall not give you any more petrol until you go.” He attempts to find out what is wrong and to set it right.⁶

Today people are often skeptical about efforts to rehabilitate criminals. Those efforts have not been notably successful, and there is reason to doubt whether they could be successful—for one thing, we do not know nearly enough about the individual causes of crime or the nature of personality or motivation to design effective ways to control them. Nevertheless, rehabilitationist ideas have been the single most important force in shaping the modern criminal justice system.⁷ In the United States prisons are not even called prisons; they are called “correctional facilities,” and the people who work in them are called “corrections officers.”

Here, then, are three theories about punishment: retributivism, deterrence, and rehabilitation. What are we to make of them? We have already seen that retributivism is more plausible than its critics suppose. But the other theories are also plausible. As for the deterrence theory, there is no doubt that sanctions are useful. They ensure massive, if imperfect, compliance with the social rules. That would probably be enough to justify punishment even if there was no other argument available. Moreover, it is hard to deny that rehabilitating criminals would be a good thing, if only we knew how to do it. Yet it can still be argued that the criminal justice system should not be designed primarily to promote deterrence or rehabilitation. Rather, it should be designed along lines suggested by the retributive theory.

⁶ Bertrand Russell, Why I Am Not a Christian (New York: Simon and Schuster, 1967), 40.

⁷ *Struggle for Justice: A Report on Crime and Punishment in America*, by the American Friends Service Committee (New York: Hill & Wang, 1971) is still one of the best discussions of the rehabilitationist idea ever produced.

The argument for this is that a system of punishment based on retributive principles is fairer and more just than systems fashioned after those other ideas. This may be shown by considering the following four principles:

Guilt: Only the guilty may be punished.

This is perhaps the most fundamental of all rules of justice: if you have committed no crime, the law should leave you alone.

Equal Treatment: People who have committed the same crime should get the same punishment.

It is not fair for one person to be sent to prison for five years, while another is incarcerated for only eighteen months, if they are charged and convicted for the same offense.

Proportionality: The punishment should be proportional to the crime.

Sometimes it is not easy to say what punishment “fits” a particular crime; nevertheless, the basic idea is clear enough. Serious crimes merit severe punishments, while minor infractions should receive only mild punishments. People should not be sent to prison for jay-walking; nor should they be fined five dollars for murder.

Excuses: People who have good excuses should not be punished, or at the very least, they should not be punished as severely as if they had no excuse.

Excuses include, for example, accident (the child ran in front of the car and there was no way the driver could stop), coercion (the man was forced to help the criminals because they were holding a gun to his head), and ignorance of fact (the nurse had been told by the child’s mother that the child was not allergic to penicillin). In each case, if there was no excuse (the driver deliberately ran over the child, the man willingly participated in the crime, the nurse knowingly gave the child a harmful drug), the person would be fully

blameworthy. But the excuse relieves the person of responsibility and so he or she should not be punished.

Any system of punishment is unjust if it departs from these four principles. But now suppose we were to design a system of punishment with deterrence in mind—that is to say, suppose we were to give the system just those features it would need to motivate people not to break the law. Would such a system satisfy these four principles, or would it violate them?

(1) *Guilt*—There is no reason, if we are concerned only with deterrence, to punish only the guilty. As far as deterrence is concerned, what matters is not whether the person punished is guilty, but whether he or she is generally *believed* to be guilty. If people believe she is guilty, the deterrent effect will be the same as if she really were guilty. Moreover, from this point of view, it would be much better to convict an innocent person (who is generally believed to be guilty) than for the crime to go “unsolved,” because when crimes are unsolved people get the idea that the law is ineffective and the deterrent effect of the law is diminished.

(2) *Equal Treatment*—A system of punishment designed solely to maximize the deterrent effect would need to be consistent in meting out similar punishments for similar crimes. This would be necessary to assure people who are tempted to violate the law that they will also get the full penalty. If in some cases lesser penalties had been imposed, then they might reasonably hope for the lesser penalty, and the deterrent effect would thereby be diminished.

(3) *Proportionality*—How severe should punishments be? If we are concerned only to deter crime, we will want to make penalties severe enough that the unlawful behavior really will be discouraged. This is a very different standard than the idea that punishments should “fit the crime.” For example, a penalty that would actually stop people from jaywalking might have to be much more severe than we would think appropriate given the trivial nature of the offense.

(4) *Excuses*—For purposes of deterrence, it is best to have a “no excuses accepted” policy. If excuses are allowed, people might hope to avoid punishment by pleading special circumstances. A system that relentlessly punishes all offenders will offer less hope of avoiding punishment and so will have much greater deterrent power.

A deterrence-based system of punishment will therefore violate three of the four principles. What about a system of rehabilitation? Rehabilitation fares no better.

(1) *Guilt*—The basic aim of such a system is to transform people who are inclined to commit crimes into people who are not inclined to commit crimes. The fact that someone *has* committed a crime is simply the best evidence we currently have of the inclination. But if it were possible to identify such people in advance, why should we wait until a crime has actually been committed? Why not go ahead and pick up individuals who are deemed likely to commit crimes and subject them to the rehabilitative routines? Of course this seems unjust, but there is nothing in the basic idea of such a system to preclude bringing people who have not committed crimes within its grasp.

(2) *Equal Treatment*—In a system designed to rehabilitate, individuals who have committed similar crimes will not receive similar treatments. What will happen to a particular lawbreaker will depend on his or her particular needs. Typically, a convicted person will be sentenced to prison for an indefinite period of time—say, “not less than ten nor more than twenty years” if the offense is serious—and then he or she will be released when the authorities (the prison officials, a parole board) decide they are “ready” to be released. Since it takes people different amounts of time to be rehabilitated, the amount of time served will vary from prisoner to prisoner.

The American criminal justice system has largely been shaped by the rehabilitationist ideal, in theory if not in fact. The widespread use of indefinite sentencing, the parole system, and the like, are manifestations of this. But the rehabilitationist character of the system has implications that are frequently misunderstood. Often, for example, critics point out that an affluent white offender is likely to serve less time in prison than a black kid from the ghetto, even if they have committed the same crime (say, a drug-related crime); and this is attributed to racism. Racism no doubt has something to do with it. But it should not be overlooked that the prevailing rehabilitationist ideology also contributes decisively to such outcomes. When the offender is well-educated, psychologically balanced, and has a good job, there’s not much for a rehabilitationist system to do with him. He may as well be given an early release. But a surly, uneducated kid with no job skills is another matter—he is just the sort of person for whom the system is designed.

(3) *Proportionality*—It should now be clear why a rehabilitative system will not respond proportionately to the crimes committed. It will respond instead to the offender's psychological or educational needs. More concretely, in most U.S. jurisdictions today, the length of one's stay in prison will depend on a parole board's judgment about when one is ready to be released, not on the seriousness of one's offense.

(4) *Excuses*—This is the only one of our four principles with which the rehabilitationist ideology is completely in accord. People need to be rehabilitated only if something in their character inclines them to commit crimes. But if someone violated the law only because of an unavoidable accident, coercion, or the like, then they do not need to be rehabilitated—or at least, the fact that they violated the law in this manner provides no evidence of it. So, in a rehabilitationist scheme, offenders with a good excuse would be let off the hook.

Once again, three of the four principles are violated. When we turn to retributivism, however, things are entirely different. Retributivism incorporates all four principles in the most natural way possible. (1) Only the guilty should be punished, because innocent people have not done anything to deserve punishment. (2) People who committed the same crime should receive the same punishment, because what one deserves depends on what one has done—it is a trivial consequence of the Principle of Desert that those who have behaved in the same way deserve the same response. (3) The Principle of Desert also requires proportionate responses, because what people deserve depends on how well, or how badly, they have behaved. A murderer has treated another person very badly indeed, and so deserves a very severe response. (That is why retributivists are inclined to support capital punishment in principle, although they might have other reasons for opposing it in practice.)

A thief, on the other hand, has done something less wicked, and so deserves a more moderate response. (4) Finally, a retributivist system of punishment would have to accept excuses, because what people deserve depends only on their *voluntary* behavior. Acceptable excuses show that behavior was not voluntary; that is why the demonstration that one was coerced, or that it was all an unavoidable accident, gets one off the hook.

We have now asked a number of questions about the three theories and four principles, and we might summarize our results like this:

	Deterrence	Rehabilitation	Retribution
Guilt	No	No	Yes
Equal Treatment	Yes	No	Yes
Proportionality	No	No	Yes
Excuses	No	Yes	Yes

The upshot is that retribution is the only idea that provides the basis of a just system of punishment. The other ideas do not even come close.

Does all this mean that Kant was right? Unlike Bentham, who believed that, to be justified, the pains of punishment must “exclude some greater evil,” Kant believed that a villain’s punishment is “good in itself, even if nothing further results from it.” The arguments we have examined seem to support Kant, but they do so only up to a point. To justify punishing someone, we may refer simply to what he or she has done—we may point out that they deserve it. But when we examine the arguments that support the general practice of treating people as they deserve, it turns out that those arguments all refer to ways in which people are *better off* under such a practice. So, at least as far as anything said here is concerned, the ultimate justifications seem to be in terms of social utility.

Thus we might understand our overall situation as follows. The best social practices are the ones that maximize welfare. The practice of treating people as they deserve is like this—people are on the whole better off if deserts are taken into account than if decisions are made solely on other grounds. One consequence of this is that we end up with a retributive understanding of punishment. Our feelings—our sense of justice, which requires that the four principles be satisfied, and our retributive feelings, which cause us to be happy when villains are punished and outraged when they are not—are useful because they support and reinforce the useful social practice. So Kant’s description of our attitudes is correct: when the annoying fellow gets at last a right good beating, we approve of it even if there are no further results. But in the larger accounting, it is a good thing that we have such attitudes only because they reinforce social practices that do have further results.