

A version of this paper appears in *Philosophical Quarterly* 63 (2013), pp. 269-292.

Two Puzzles About Mercy

Ned Markosian

Abstract: Anselm raised a puzzle about mercy: How can anyone (God, say, or a judge) be both just and merciful at the same time? For it seemed to Anselm that justice requires giving people what they deserve, while being merciful involves treating people less harshly than they deserve. This puzzle has led to a number of analyses of mercy. But a strange thing emerges from discussions of this topic: people seem to have wildly divergent intuitions about putative cases of mercy. Examples that are taken by some to be paradigm cases of mercy are seen by others to be clear instances of non-mercy. This phenomenon gives rise to a new puzzle: Why is it that people seem to have such wildly divergent intuitions about putative examples of mercy? This paper proposes a new analysis of mercy that provides solutions to both of these puzzles about mercy.

1 Introduction

There is an old puzzle about justice and mercy. It was first described in print by Anselm, who wondered how it could be that God is both just and merciful.¹ For it seemed to Anselm that to be just is to give people what they deserve, and that being merciful involves treating people less harshly than they deserve. (Although the thesis that justice is done when people get what they deserve is both ancient and highly plausible, it has nevertheless been denied by some. For the purposes of this paper, however, I will simply assume that thesis without argument.) This puzzle about justice and mercy is of interest to atheists as well as theists, since the difficulty is a general one: How can anyone – a judge, for example – be both just and merciful? Largely as a result of consideration of Anselm's puzzle, many philosophers have

¹ Anselm, *Proslogium*, Ch. IX. Anselm also discussed another puzzle about mercy:

attempted to give a correct analysis of the concept of mercy. Some have appealed to their own proposed analyses of mercy in attempts to show that a single act can be both just and merciful, while others have proposed analyses of mercy and then appealed to those analyses in attempts to show that no act can have both of these features.

Apart from its relevance to Anselm's puzzle about justice and mercy, the question of what is the correct analysis of the concept of mercy is interesting in its own right. So it seems appropriate that the question has received a fair amount of attention. But a strange thing emerges from discussions of this question: it turns out that people have wildly divergent intuitions about certain putative cases of mercy. Examples that are taken by some to be paradigm cases of mercy are seen by others to be clear instances of non-mercy, in a way that goes far beyond the normal failure of philosophers to see eye-to-eye on specific cases.² This phenomenon gives rise to a new puzzle about mercy: Why is it that people seem to have such wildly divergent intuitions about putative examples of mercy? Is there a single concept of mercy that all or most of us have in mind when we use words like 'mercy' and 'merciful'? If so, then why do people disagree so much on the examples? If not, then what are we all talking about when we use these words?

The aim of this paper is to propose a new analysis of the concept of mercy that provides solutions to both of the puzzles about mercy. First I will make some preliminary remarks that will focus the ensuing discussion of analyses of mercy. Next I will look at the most plausible analyses of mercy that have been suggested in the literature (together with several other plausible accounts that have not yet shown up in the literature), considering, in the case

² There are many instances of this phenomenon in the literature, where it is common to find one author claiming that a certain example is a clear case of mercy and another author citing the same example as an obvious case of non-mercy. (Rainbolt's Nazi doctor (see below) is a case in point – Rainbolt thinks it is a clear case of mercy, while I think it is just the opposite. Other cases are mentioned below.) To get an even better appreciation of the phenomenon, one has merely to participate in informal discussions with others of various putative cases of mercy. This often leads to incredulous exclamations along the lines of "You think *that* is a case of mercy? You must be crazy!" and "But it's completely obvious that this is a classic case of mercy!" In my experience, the amount of disagreement over cases in such conversations is almost comical, and seems to be far greater than is normal in discussions among philosophers.

of each such account, what appear to be serious problems for that analysis. Finally, I will present and defend a new analysis of mercy, and attempt to show that it provides satisfactory solutions to both Anslem's puzzle about justice and mercy and the new puzzle about wildly divergent intuitions concerning specific putative examples of mercy.

2 Prolegomena to Any Future Analysis of Mercy

I take it that the adjective 'merciful' can be appropriately applied to both people and actions. But I think that the mercy exhibited by actions is the more fundamental variety, so I will here be concerned with analyses of the concept of a merciful action. (After all, it is natural to say that a merciful person is just a person who is inclined to perform merciful actions.) I also take it that each merciful action involves not only its agent but also a recipient, i.e., the one to whom the agent is being merciful. As I see it, then, the fundamental locution concerning mercy is something like this: ' x is merciful to y at t ', where the terms in place of ' x ' and ' y ' refer to agents, and the term in place of ' t ' refers to a time.

It will be useful to have some uncontroversial cases of mercy before us. Here are what I take to be two such cases:

Case 1: A judge is sentencing a convicted thief. The law calls for anywhere from 5 to 10 years in prison. The judge, moved by compassion for the thief, sentences her to 5 years.

Case 2: A soldier has defeated another soldier in battle. The victorious soldier holds a sword to the throat of her foe, who has been disarmed. But the defeated soldier pleads for mercy, and the victorious soldier, moved by compassion for her vanquished foe, releases him.

It will also be useful to have at our disposal a test that can help to determine, in the case of any given action, whether that action is a merciful one. I propose the following test. Suppose that we want to decide whether x was merciful to y at some time, t . Imagine that we have reunited x and y after the fact, in a situation in which x must plead his or her case before y , who is to be the sole arbiter in determining x 's fate. Suppose that x is allowed to appeal to various facts of the case in an attempt to convince y that x was merciful to y at t . But suppose that, as a result of taking some powerful truth serum, it is

physically impossible for x to say anything false in pleading his or her case. Imagine that if x can convince y that x was merciful to y at t , then x will receive a favorable decision from y ; otherwise, x will receive an unfavorable decision. I will refer to this imaginary situation, to be considered in attempts to determine whether a given action is merciful, as “Mercy Court.”³

It is clear that in order to have any chance of winning a favorable decision in Mercy Court, x will have to be able to say that he or she *could* have treated y more harshly at t than x in fact did. And it is also clear that x will have to be able to say that he or she chose the less harsh treatment out of compassion for y (and not as a result of purely selfish considerations, or out of concern for some third party,⁴ or solely from a belief that any harsher treatment would have been in some way impermissible).

What remains unclear, however, is what is the relevant sense of ‘could’ in the sentence ‘ x could have treated y more harshly at t than x in fact did’. It cannot be the *could* of logical possibility, for if it were then virtually every compassionate act would count as merciful. Likewise, it cannot be the *could* of physical possibility, for similar reasons. (And for an additional reason: most of us want the performance of merciful acts to be compatible with determinism about the laws of nature.) So there must be some interesting way of answering the question *What is the sense of ‘could’ in the sentence ‘ x could have treated y more harshly at t than x in fact did’ that is relevant to determining whether x was merciful to y at t ?* And it seems to me that what distinguishes each plausible analysis of mercy (including every one that has been suggested in the literature so far, as well as the several others discussed below that haven’t been suggested in the literature) from its rivals is its peculiar way of answering this question.

The way I see it, then, each plausible analysis of mercy can be formulated as an instance of the following schema.

³ A useful variation on Mercy Court: It is not the putative recipient of mercy (y) who sits in judgement of the agent in question (x). Instead, it is you and me. But the setup is the same: the agent whose action is under consideration as a possible act of mercy must plead his or her case, and benefits from a favorable ruling, but is only permitted to say things that are true.

⁴ If x admits to choosing the the less harsh treatment for y out of concern for a third party, then y can legitimately insist that x ’s action may have been an instance of mercy toward that third party, but is not in any case an instance of mercy toward y .

A Schema for Analyses of Mercy: x is *merciful* to y at t =df there are two alternative acts, A and B , available to x at t such that (i) x performs A , (ii) x 's performing B would have resulted in more harm to y than x 's performing A , (iii) x is motivated by compassion for y to choose A over B , and (iv)

where the fourth condition represents a way of spelling out the relevant sense of 'could' in the sentence ' x could have treated y more harshly at t than x in fact did'.

3 Some Plausible Analyses of Mercy

Let us begin our consideration of analyses of mercy with one that, while initially very plausible, has not, to my knowledge, been defended in the literature. According to this analysis, the answer to the above question about the relevant sense of 'could' is that it is a *prudential could*. That is, when x says to y in Mercy Court, "I could have treated you more harshly than I in fact did," x must mean that it would have been prudent for x to perform the action that would have resulted in more harm to y . This analysis of mercy can be spelled out as follows.

The Prudential Analysis of Mercy (PAM): x is *merciful* to y at t =df there are two alternative acts, A and B , available to x at t such that (i) x performs A , (ii) x 's performing B would have resulted in more harm to y than x 's performing A , (iii) x is motivated by compassion for y to choose A over B , and (iv) x 's performing B would have been at least as prudent as x 's performing A .

Although PAM is initially plausible, I think it is pretty clear that it yields some unacceptable results. Consider the following case.

Case 3: Smith and Jones are both in line for the same promotion, but Smith has learned that the brass is leaning toward giving the job to Jones. One day, Smith has an opportunity to kill Jones, who happens to be an innocent man and an upstanding citizen, from several hundred yards away using a high-powered rifle. Smith knows that she would get away with this murder if she were to commit it, and also that

things would work out very well for her if Jones were out of the picture. She also knows that she would feel no compunction about doing the deed, nor any subsequent guilt. After sighting Jones through the scope of her rifle, Smith feels a wave of compassion for the poor man, and decides not to kill him, even though she knows that she could get away with it, and would be better off if she did.

Since killing Jones would be more prudent for Smith than not killing him, and since Smith chooses not to kill him out of compassion for Jones, PAM entails that Smith's not killing Jones is a case of mercy. But it seems clear that this is not at all an act of mercy: Jones would have to be crazy to hand down a favorable Mercy Court ruling concerning Smith's actions in this case.

Consideration of PAM and this sort of bad result that it generates may suggest a different analysis of mercy. This second analysis is hinted at by Anselm in the passage in which he raises his puzzle about justice and mercy,⁵ and is one that I think many people would be inclined to accept.⁶ The idea behind this analysis is that the relevant sense of 'could' in *x*'s Mercy Court claim that *x* could have treated *y* more harshly has to do with the notion of desert. According to this line of thinking, *x* is being merciful to *y* only if there is an alternative available to *x* that involves treating *y* more harshly, but in a way that *y* actually deserves (or deserves as much as the relevant lenient treatment). This idea can be captured in the following analysis of mercy.

The Desert Analysis of Mercy (DAM): *x* is *merciful* to *y* at *t* =df there are two alternative acts, *A* and *B*, available to *x* at *t* such that (i) *x* performs *A*, (ii) *x*'s performing *B* would have resulted in more harm to *y* than *x*'s performing *A*, (iii) *x* is motivated by compassion for *y* to choose *A* over *B*, and (iv) the harsher treatment that *y* would have received as a result of *x*'s performing *B* is at least as close to what *y* deserves as the less harsh treatment that *y* receives as a result of *x*'s performing *A*.

⁵ Anselm, *Proslogium*, Ch. IX.

⁶ See for example H.M Hurd, 'The Morality of Mercy', *Ohio State Journal of Criminal Law* 4 (2007), pp. 389-421.

Unfortunately, however, DAM is open to the following objection. It entails that if the vanquished soldier in Case 2 above deserves to live more than he deserves to die, then the victorious soldier's act of releasing him is not an act of mercy. It seems to me that this is an unacceptable consequence. I think that, according to the way we ordinarily use the word 'merciful', the victorious soldier's action should count as a merciful action, even if the vanquished soldier does deserve to live more than he deserves to die.⁷ (Supporters of DAM may disagree with me here, and insist that the victorious soldier's action in Case 2 is not really an example of mercy. I will have more to say about this disagreement below in the discussion of my own proposed analysis of mercy.)

It might be thought that the proponent of DAM could get around the above objection by insisting that the vanquished soldier in Case 2 does in fact deserve to die. For it might be argued that a soldier who willingly goes into battle has implicitly accepted the rules of warfare, one of which is that if you fight and lose then you deserve to die. I myself do not find this line of reasoning at all plausible. But in any case, the example can be qualified so that this response is avoided. We merely have to stipulate that the vanquished soldier did not go into battle willingly but, rather, was forced or somehow coerced into fighting. If the soldier is thus made into an innocent victim of cruel but higher powers, then the idea that he deserves to die for being defeated in battle loses whatever plausibility it initially had.

H. Scott Hestevold has proposed a variation on the Desert Analysis of Mercy.⁸ Hestevold's analysis involves the notion of "disjunctive deserts," or "disjunctions of penalties, each disjunct of which could alone serve as the sufficient desert for the improper act in question."⁹ Hestevold claims that in such cases, "the sufficient deserts may be of varying severity."¹⁰ Hestevold's analysis of mercy is based on the assumption that there can be such

⁷ Rainbolt raises a similar objection to a proposal relevantly like DAM in G. Rainbolt, 'Mercy: In Defense of Caprice', *Nous* 31 (1997), pp. 226-241, at pp. 228-229.

⁸ In H.S. Hestevold, 'Disjunctive Desert', *American Philosophical Quarterly* 20 (1983) 357-363; and H.S. Hestevold, 'Justice to Mercy', *Philosophy and Phenomenological Research* 46 (1985), pp. 281-291.

⁹ See 'Justice to Mercy', p. 281, and also 'Disjunctive Desert', p. 360.

¹⁰ 'Justice to Mercy', p. 281.

disjunctive deserts. The idea behind Hestevold's analysis is that the relevant sense of 'could' in the sentence '*x* could have treated *y* more harshly at *t* than *x* in fact did' involves this notion of disjunctive deserts. If there are several alternatives available to *x*, each of which involves giving *y* a member of a disjunctive desert that *y* deserves, and if those disjuncts are of varying severity, then *x* is in a position to be merciful toward *y*. Although he does not formulate his analysis in the terms we have been employing, Hestevold's analysis can, for our purposes, be formulated as follows.

Hestevold's Analysis of Mercy (HAM): *x* is *merciful* to *y* at *t* =df there are two alternative acts, *A* and *B*, available to *x* at *t* such that (i) *x* performs *A*, (ii) *x*'s performing *B* would have resulted in more harm to *y* than *x*'s performing *A*, (iii) *x* is motivated by compassion for *y* to choose *A* over *B*, and (iv) *x*'s performing *A* results in *y*'s receiving one disjunct of a disjunctive desert that *y* deserves, while *x*'s performing *B* would have resulted in *y*'s receiving another disjunct of that disjunctive desert.¹¹

This proposal provides an ingenious solution to Anselm's puzzle about justice and mercy. For if HAM is correct, then there can be cases in which a person is both just and merciful. Indeed, if HAM is correct, then every case of mercy is also a case of justice. (Recall that I am assuming the somewhat controversial thesis that justice is done when people get what they deserve.) For in order for *x* to be merciful to *y*, according to HAM, *x* must give *y* one of the sufficient deserts from among the disjuncts of a disjunctive desert.

Still, there are at least two serious objections to HAM. The first is that it entails that *x* cannot be merciful to *y* unless there is some disjunctive desert of penalties that *y* deserves. But suppose that the thief in Case 1 above deserves to be sentenced to 10 years in prison, and does not deserve any other sentence. It seems to me that the judge would then be acting mercifully toward the thief if, as a result of compassion for the thief, he sentenced her to 5 years rather than 10 years. (And I maintain that most people will agree with me that the judge in this case would win a favorable decision from the thief in Mercy Court.) But according to HAM, the judge's act in this case would not count as

¹¹ See Hestevold, 'Disjunctive Desert', pp. 360-61, and Hestevold, 'Justice to Mercy', p. 281.

merciful, since the penalty in question (5 years) is not a member of any relevant disjunctive desert.

A second objection to HAM is that it is open to the same criticism that was raised above in connection with DAM: it entails that if the vanquished soldier in Case 2 above does not deserve to die, then the victorious soldier does not act mercifully when she releases her foe. This seems to me to be an unacceptable result.

I suspect that these objections would not convince Hestevold to reject HAM. I suspect that his response to the first objection would be to insist that the relevant version of Case 1 (without the presence of a disjunctive desert) is not a case of mercy. And I suspect that his response to the second objection would be (similarly) to maintain that the victorious soldier's act in Case 2 is not an example of mercy (on the assumption that it would be unjust for her to kill her foe, that is). Moreover, I suspect that Hestevold would find this consequence perfectly intuitive. Thus, I think that we have here two further examples in which the intuitions of different people (in this case, Hestevold and myself) regarding specific putative cases of mercy diverge wildly.

In any event, Alwynne Smart has suggested an alternative analysis of mercy.¹² She says the following.

When a man exercises mercy, what he does is acknowledge that an offense has been committed, decides that a particular punishment would be appropriate or just, and then decides to exact a punishment of lesser severity than the appropriate or just one.¹³

This passage suggests an analysis of mercy according to which the relevant sense of 'could' in the sentence '*x* could have treated *y* more harshly at *t* than *x* in fact did' has to do with justice. The idea is that if it would have been just for *x* to punish *y* more harshly than *x* in fact does, then *x*'s action may be merciful. This idea can be captured in the following analysis of mercy.

The Justice Analysis of Mercy (JAM): *x* is *merciful* to *y* at *t* =df there are two alternative acts, *A* and *B*, available to *x* at *t* such that (i) *x* performs *A*, (ii) *x*'s performing *B* would have resulted

¹² A. Smart, 'Mercy', *Philosophy* 43 (1968), pp. 345-359.

¹³ 'Mercy', p. 350.

in more harm to y than x 's performing A , (iii) x is motivated by compassion for y to choose A over B , and (iv) it would have been just for x to perform B .

JAM will no doubt strike many people as plausible, but it should be clear that, at least given our assumption that justice is done when people get what they deserve, JAM is open to the same objection raised against DAM and HAM: it entails that if the vanquished soldier in Case 2 above does not deserve to be killed, then the victorious soldier's action in that case cannot count as an example of mercy. And as I have said, this seems like an unacceptable result.

I take it that this objection would not convince Smart to reject JAM, and that her response would be to maintain that the victorious soldier's act in Case 2 is not an example of mercy, on the assumption that it would be unjust for her to kill her foe. And I suspect that Smart would find this consequence perfectly intuitive. Thus, I think that we have here yet another example in which the intuitions of different people (in this case, Smart and myself) regarding a specific putative case of mercy diverge wildly. I will have more to say below about how the different intuitions can be resolved in this particular case.

What's more, it should also be clear that JAM leaves us with no way of resolving Anslem's puzzle about justice and mercy. For if it would have been just for x to perform B , and if B is worse for y than A , then it is hard to see how it could be just for x to perform A . Thus JAM seems to entail that it is not possible for anyone to be both just and merciful at the same time. Insofar as many people think that justice and mercy *are* compossible, this appears to be a major strike against JAM.

John Kleinig has suggested yet another analysis of mercy. He writes,

When we look at mercy in this broader context we can see that its core idea is that of treating with benevolence those who are in need, distress, debt or under threat of some sort, and who have no claim to such benevolent treatment.¹⁴

The view suggested by Kleinig is that the sentence ' x could have treated y more harshly at t than x in fact did' is to be understood as meaning that y has

¹⁴ J. Kleinig, 'Mercy and Justice', *Philosophy* 44 (1969), pp. 341-342, at p. 341.

no claim to the less harsh treatment that y receives from x . This view can be formulated as follows.

The Claims Analysis of Mercy (CAM): x is *merciful* to y at t =df there are two alternative acts, A and B , available to x at t such that (i) x performs A , (ii) x 's performing B would have resulted in more harm to y than x 's performing A , (iii) x is motivated by compassion for y to choose A over B , and (iv) y has no claim to the less harsh treatment that results from x 's performing A .

CAM is initially very appealing, but it is open to at least one serious objection. Consider the following case.

Case 4: Brown has the responsibility of sentencing a man convicted of a crime. But the convicted man, Red, has paid Brown a considerable sum of money as a bribe. The agreement between Brown and Red is that in exchange for the bribe, Brown will give Red the lightest possible sentence allowed by the law. Before the sentencing, however, Brown has an attack of conscience, and decides never to take another bribe. She also decides that she will give the money from Red to a charity, and will ignore the fact that he tried to bribe her. Then, moved by compassion for Red, Brown sentences him to the lightest possible sentence anyway.

CAM entails that Brown is not merciful to Red in this case, since Red *does* have a claim to the less harsh treatment that results from Brown's handing down the lighter sentence. (Imagine Red saying to his cronies, "Well, sure I got a light sentence, but it was bought and paid for, wasn't it?") But it is plausible to think that Brown's action in Case 4 is indeed a case of mercy. For Brown could make the following speech in Mercy Court.

I realize that you thought you had effectively bribed me, Red, but as it turned out you had not done so. I had an attack of conscience, you see, and decided not to be bribed ever again, not by you or by anyone else. Therefore, whether fairly or unfairly, I decided to ignore the fact that you had a claim to the lighter sentence. Instead, at the time of my decision, all I took into account were the different possible sentences available to me, and the feeling of deep compassion that I felt for you. I

knew that I could give you a harsher sentence, but I wanted to be nice to you, so I gave you the lighter sentence.

It seems to me that Red, upon hearing this speech, and realizing that everything Brown says in Mercy Court must be true, would surely agree that Brown had shown him mercy.¹⁵

Another case that seems to be problematic for CAM is one in which y has some claim to the less harsh treatment that results from x 's performing A , but x , acting out of compassion for y , and *without any knowledge of the fact that y has a claim to the less harsh treatment*, chooses A over B .

Consideration of certain standard examples of mercy, such as Case 1 above, in which the judge sentences the thief to 5 years rather than 10, suggests that the sense of 'could' that is relevant to mercy is a *legal could*. Here is an analysis of mercy based on this idea.

The Legal Analysis of Mercy (LAM): x is *merciful* to y at t =df there are two alternative acts, A and B , available to x at t such that (i) x performs A , (ii) x 's performing B would have resulted in more harm to y than x 's performing A , (iii) x is motivated by compassion for y to choose A over B , and (iv) B would have been permitted by the laws governing the situation.¹⁶

Although LAM is relatively plausible, it faces what I think are insurmountable problems. For one thing, suppose the state passes a law tomorrow making it legal for a person to kill anyone else at any time. Then there will be outbreaks of mercy all over the place, according to LAM, although that result flies in the face of common sense. Also, consider the situation in a "state of nature," in which there are no laws at all. Then any act that involves compassionately refraining from killing another to steal his dinner will count as an act of mercy, according to LAM; but surely that is an absurd consequence.

¹⁵ But an anonymous reader has informed me that she or he feels strongly that Case 4 is *not* an instance of mercy, so that CAM actually gets the right result here. Apparently Case 4, then, is another putative example of mercy concerning which different people seem to have surprisingly divergent intuitions.

¹⁶ Dolinko seems to have in mind something like LAM in D. Dolinko, 'Some Naïve Thoughts About Justice and Mercy', *Ohio State Journal of Criminal Law* 4 (2007), pp. 349-360.

George Rainbolt has suggested an account of mercy that many will consider to be an improvement on LAM.¹⁷ Here is (my formulation in our terms of) Rainbolt's analysis.

The Normal Analysis of Mercy (NAM): *x* is *merciful* to *y* at *t* =df there are two alternative acts, *A* and *B*, available to *x* at *t* such that (i) *x* performs *A*, (ii) *x*'s performing *B* would have resulted in more harm to *y* than *x*'s performing *A*, (iii) *x* is motivated by compassion for *y* to choose *A* over *B*, and (iv) *B* would have been normal for its circumstances.

NAM seems to get better results than LAM in a wide variety of cases. But there are nevertheless problems facing NAM. Consider a case that Rainbolt himself introduces: There is a Nazi doctor whose job is to torture innocent children. But there is something about one of the children that strikes a chord in the doctor. As a result of the compassion that he feels for this one child, the doctor chooses not to torture her. Is this an act of mercy? According to NAM it is – for torturing the child would be the normal thing for the Nazi doctor to do, under the circumstances. But it seems to me that this is a very counterintuitive result, for the simple reason that it would be wrong of the doctor to torture the child in the example. (It's true that torturing the child would be just one more immoral act in a series of such acts, but that doesn't seem to make it a case of mercy when the doctor chooses not to torture the child.)¹⁸

¹⁷ G. Rainbolt, 'Mercy: An Independent, Imperfect Virtue', *American Philosophical Quarterly* 27 (1990), pp. 169-173. See also his 'Mercy: In Defense of Caprice', where Rainbolt defends a different analysis of mercy, according to which (in our terminology) it must be the case that *x* has a "fairly strong reason" to perform *B*. I think that remarks similar to those in the text about NAM would apply to Rainbolt's new variation on NAM.

¹⁸ Perhaps a proponent of NAM who agrees with me (and disagrees with Rainbolt) that the Nazi doctor's act is not a case of mercy could argue that NAM gets the right result in this case because torturing children is not normal. I am sympathetic to this response. But I think that the best argument for the claim that, even in the Nazi doctor's situation, torturing children is not normal would be an argument that appeals either to the idea that torturing children is not *just* or to the idea that torturing children is not *moral*. And if a proponent of NAM were to defend NAM by appealing

Interestingly enough, Rainbolt introduces the example of the Nazi doctor because he thinks that it clearly *is* a case of mercy. That strikes me as extremely counterintuitive – but I am not very surprised that Rainbolt (and others, no doubt, who share his intuitions about the case) should have intuitions about an alleged case of mercy that are so radically different from my own. For as we have already noted, one of the puzzling things about mercy is the fact that different people tend to have wildly divergent intuitions about specific cases. I will have more to say below about the Nazi doctor case, and why Rainbolt and I should have such diametrically opposed reactions to that case.

Consideration of the case of the Nazi doctor suggests that the *could* in our sentence ‘*x* could have treated *y* more harshly at *t* than *x* in fact did’ should be understood as a *moral could*. This line of thinking gives us an account of mercy that strikes me as very plausible, although no one has yet suggested it in the literature. Here is the account.

The Moral Analysis of Mercy (MAM): *x* is *merciful* to *y* at *t* =df there are two alternative acts, *A* and *B*, available to *x* at *t* such that (i) *x* performs *A*, (ii) *x*’s performing *B* would have resulted in more harm to *y* than *x*’s performing *A*, (iii) *x* is motivated by compassion for *y* to choose *A* over *B*, and (iv) *B* would have been morally permissible.

It seems to me that MAM gives correct results in a wide variety of cases, including the case of the Nazi doctor. (For, given that torturing the child would not have been morally permissible, MAM implies that the Nazi doctor’s sparing of the child is not in fact an act of mercy.) But there are at least three major problems facing MAM. First, depending on which ethical theory happens to be true, it will be possible to construct counterexamples along the lines of Case 2 above. For example, suppose that act utilitarianism is true, and that killing the vanquished soldier would not maximize utility. Or suppose that Rossian Minimalism is true, and that killing the vanquished soldier would fail to minimize prima facie duty violations.¹⁹ Then, on either of

to either of these two ideas, then I think he or she would be tacitly committed to either JAM or MAM (see below).

¹⁹ On Rossian Minimalism see N. Markosian, ‘Rossian Minimalism’, *Journal of Ethics and Social Philosophy* 4 (2009), pp. 1-16. I mention act utilitarianism and Rossian Minimalism here merely to provide concrete examples, but I hope that readers who

these suppositions, it would be morally wrong for the victorious soldier to kill her foe, which means that according to MAM, Case 2 will not be a case of mercy. And this seems to be the wrong result.

Second, many people believe that it can be morally obligatory, in certain situations, for an agent to be merciful. (For example, it is a widely held view that Shylock, in Shakespeare's *The Merchant of Venice*, is morally obligated to be merciful toward Antonio by forgoing his pound of flesh.) But if this is right, then MAM must be wrong, since MAM, in virtue of its fourth condition, entails that mercy can never be morally obligatory.²⁰

What's more, even if MAM could somehow get around these two difficulties, it would still have another shortcoming: it doesn't solve the second puzzle about mercy, the one about the wildly divergent intuitions of different people. To solve that puzzle, we will need a more flexible analysis of mercy.

Before turning to such an account, there is one more proposal that a number of authors have made, and that I should deal with at least briefly. Claudia Card, Jeffrie Murphy, and P. Twambley have all (independently) suggested an analysis of mercy that can be formulated in the following way.²¹

The Rights Analysis of Mercy (RAM): *x* is *merciful* to *y* at *t* =df there are two alternative acts, *A* and *B*, available to *x* at *t* such that (i) *x* performs *A*, (ii) *x*'s performing *B* would have resulted in more harm to *y* than *x*'s performing *A*, (iii) *x* is motivated by compassion for *y* to choose *A* over *B*, and (iv) *x* had a right to perform *B*.

Here is the main problem facing RAM. What kind of right is involved in the analysis? A legal right? If so, then RAM is equivalent to LAM, and is open to the same objections as that view. Or is the kind of right involved in RAM a

endorse some other ethical theory will be able to agree that there are many plausible ethical theories that yield the same result, namely, that it would be morally wrong for the victorious soldier in Case 2 to kill her vanquished foe.

²⁰ I am grateful to an anonymous referee for bringing up this objection to MAM.

²¹ See J. Murphy, and J. Hampton, *Forgiveness and Mercy* (Cambridge University Press, 1988), p. 176; P. Twambley, 'Mercy and Forgiveness', *Analysis* 36 (1975), pp. 84-90, at p. 86; and C. Card, 'On Mercy', *Philosophical Review* 81 (1972), pp. 182-207, at p. 182.

moral right? If it is, then the view is equivalent to MAM, and is open to the objections raised against that view.

4 The Flexible Analysis of Mercy

Our second puzzle about mercy concerned the fact that people seem to have wildly divergent intuitions about putative examples of mercy. This phenomenon calls for some explanation. One possible explanation is that there simply is no single concept of mercy that is common to all, or even most, people. Thus, according to this line, when people seem to be disagreeing about whether a particular case is an example of “mercy” they are not really disagreeing at all; rather, they are talking about completely different concepts. One party is saying that the case is an example of one concept, and the other party is saying that the case is not an example of some *other* concept. They are simply talking past each other.

The difficulty with this explanation for the seemingly wildly divergent intuitions is that it is just too depressing. It would be a shame if we had to admit that there is no common concept of mercy, and that we are generally talking past each other, or making meaningless sounds, when we use words like ‘mercy’ and ‘merciful’. It would be much better to have an explanation of the relevant phenomenon according to which there is a communal concept of mercy that is grasped by lots of different people. The analysis of mercy that I want to defend allows us to explain the phenomenon in just such a way. For on the analysis that I am proposing, there is a single, communal concept of mercy, but it is one that is flexible in an important way. On the analysis I am proposing, the force of ‘could’ in ‘*x could have done the harsher thing*’ is something like this: *x’s doing the harsher thing would have been permitted under the relevant notion of permissibility*. And the notion of permissibility that is relevant to determining whether *x* could have done otherwise, on the view I want to propose, can vary from one context to another. How this will allow us to explain the phenomenon in question (of wildly diverging intuitions about specific cases), while still maintaining that there is a single, communal concept of mercy will, I hope, become clear below. But first let me explain the analysis of mercy that I am proposing.

The analysis I am proposing is inspired by the anti-fatalist, “contextualist” analysis of ability spelled out by David Lewis.²² According to Lewis’s analysis of ability, to say that x is *able* to do A is to say that x ’s doing A is consistent with the relevant facts. And it is an important part of Lewis’s proposal that which facts are relevant in a given case is determined by features of the context of utterance. So, for example, if you and I are physiologists discussing whether Lewis can speak Finnish, then the main facts that will be relevant are facts about the anatomy and operation of Lewis’s larynx and nervous system; so that if you say that Lewis can speak Finnish, what you say will be true. But if Tim and Tom are language instructors discussing whether Lewis can speak Finnish, then the main facts that will be relevant are facts about Lewis’s upbringing and training; so that if Tom says that Lewis *cannot* speak Finnish, then what Tom says will also be true. Another important part of Lewis’s proposal is the idea that in some contexts, the matter of which facts are relevant to the question of whether a certain person can perform a certain action will be underdetermined.

The analysis of mercy that I want to propose – which I will call the “Flexible Analysis of Mercy” – is based on Lewis’s analysis of ability. Here is a preliminary formulation of the view.

The Flexible Analysis of Mercy, Preliminary Version (FAM-PV): x is *merciful* to y at t =df there are two alternative acts, A and B , available to x at t such that (i) x performs A , (ii) x ’s performing B would have resulted in more harm to y than x ’s performing A , (iii) x is motivated by compassion for y to choose A over B , and (iv) x ’s performing B would have been permitted under the relevant notion of permissibility.

²² See D. Lewis, ‘The Paradoxes of Time Travel’, *American Philosophical Quarterly* 13 (1976), pp. 145-152. More recently there has been an outbreak of similar, contextualist views in epistemology. See for example S. Cohen, ‘Knowledge and Context’, *The Journal of Philosophy* 83 (1986), pp. 574-583; S. Cohen, ‘Knowledge, Context, and Social Standards’, *Synthese* 73 (1987), pp. 3-26; S. Cohen, ‘How to be a Fallibilist’, *Philosophical Perspectives, Volume 2* (1988), pp. 91-123; S. Cohen, ‘Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery’, *Australasian Journal of Philosophy* 76 (1998), pp. 289-306; and K. DeRose, ‘Contextualism and Knowledge Attributions’, *Philosophy and Phenomenological Research*, 52 (1992), pp. 913-929.

The idea here is that which notion of permissibility is relevant in a given case is determined (and sometimes underdetermined) by features of the context of utterance. So, for example, if some lawyers are discussing whether the judge in Case 1, who was moved by compassion for a thief to sentence her to 5 years instead of 10, acted mercifully, then the relevant notion of permissibility will no doubt be *legal permissibility*. In particular, since sentencing the thief to 10 years would have been consistent with the laws governing the case, condition (iv) of the above analysis will be satisfied, and the case will count as an example of mercy.²³

FAM-PV captures the basic idea behind the analysis I want to defend, but an amendment may be needed in order to avoid a particular kind of counterexample.²⁴ Here is an instance of the type of counterexample I have in mind.

Case 5: Alice and Betty are leading physicists and candidates for the Nobel Prize. It just so happens that Alice will get the prize if she lives, but Betty will get it otherwise. It also just so happens that things will be very bad for Betty if she doesn't get the prize, and much less bad if she does get it. So, out of compassion for Betty, her sister, Carla, kills Alice, thereby ensuring that Betty will get the prize.

Here's why Case 5 is a counterexample to FAM-PV. In normal contexts it will be true to say that Carla's killing Alice counts as an instance of mercy toward Betty, according to FAM-PV, since the case will satisfy all four conditions of the analysis. In particular, in normal contexts it will be true to say that Carla's *not* killing Alice would have been permitted under the relevant notion of permissibility, given any plausible account of how the relevant notion of permissibility is determined. But it seems pretty clear that this is not at all a case of mercy – it is instead a case of a cruel and gratuitous murder of an

²³ The reader might wonder whether I have a complete theory of how it is determined, in a given context, which notion of permissibility is relevant to claims about mercy. Unfortunately I do not. (One can only do so much in a single paper.) But it is worth pointing out that a similar question arises in connection with contextualist theories in epistemology. (See the works by Cohen and DeRose cited in the previous footnote.)

²⁴ I am indebted here to Frances Howard-Snyder.

innocent, Nobel-Prize-worthy physicist. (Notice that variations on Case 5 will also constitute counterexamples to most of the other analyses of mercy considered so far in this paper. Notice, also, that in each case there will be an amendment along the lines of the amendment to FAM-PV spelled out below that will get around the relevant counterexample.)

Although FAM-PV is susceptible to counterexamples of this kind, I think that The Flexible Analysis of Mercy can be modified so that it avoids these counterexamples. What is needed is a simple change to the fourth condition. Not only *B*, the harsher alternative, but also *A*, the candidate for being a merciful action that was in fact performed, must be permitted under the relevant notion of permissibility. Thus the new formulation of the analysis will look like this.

The Flexible Analysis of Mercy (FAM): *x* is *merciful* to *y* at *t* =df there are two alternative acts, *A* and *B*, available to *x* at *t* such that (i) *x* performs *A*, (ii) *x*'s performing *B* would have resulted in more harm to *y* than *x*'s performing *A*, (iii) *x* is motivated by compassion for *y* to choose *A* over *B*, and (iv) *A* and *B* are both permitted under the relevant notion of permissibility.

In order to get a better understanding of how FAM works, let's take another look at some of the examples from above to see what kind of results FAM will generate. In most contexts in which Case 1 (involving a judge who sentences a thief to 5 years rather than 10) is under discussion, the notion of *legal permissibility* would no doubt be the relevant one, so that FAM implies that it would be true to say in those contexts that the judge's action is a case of mercy. Similarly, in discussions of Case 2 (the case involving the soldiers), the notion of *wartime permissibility* would presumably be the relevant notion in normal contexts, so that FAM implies that in normal contexts it would be true to say that the victorious soldier's action is a case of mercy.²⁵ And likewise, in discussing Case 3 (the one involving Smith, who chooses not to kill Jones, her rival for a big promotion), most people would no doubt take the notion of *moral permissibility* to be the relevant one in nearly every context, so that FAM implies that it would nearly always be true to say that Smith's action is *not* a case of mercy.

²⁵ But see the third objection to FAM discussed in the final section of this paper.

Another important consequence of FAM is that the phenomenon of underdetermination will play an important role in our talk about mercy. For it is likely that quite often the features of a given context of utterance will underdetermine the matter of which notion of permissibility is the relevant one with respect to a certain putative case of mercy. When that happens, according to FAM, there will simply be no fact of the matter, relative to those underdetermined contexts, about whether the case is one of mercy. If this is right, it would go a long way toward explaining why it is so hard to decide whether mercy is involved in many cases, and also why people disagree so much about specific putative cases of mercy. These phenomena are exactly what we should expect in an area affected by such underdetermination.

Note, however, that this is not to say that there is no single concept of mercy that the different people are talking about when they disagree over a specific case. Nor is it to say that there is no single concept of mercy that is relevant to the different controversial cases. On the contrary, there is a single concept of mercy that is common to different people and different cases, according to FAM. It is exactly the concept captured by FAM.

5 FAM and the Two Puzzles About Mercy

I claimed at the outset that my proposed analysis of mercy provides solutions to both of our puzzles about mercy. Let me sketch out the solutions I have in mind. First, consider Anselm's puzzle about justice and mercy. How can someone be both just and merciful at the same time? Well, it's easy. Take the judge in Case 1. As long as his action satisfies conditions (i)-(iii) in FAM, and as long as we are in a context in which, say, the legal notion of permissibility is the relevant one, then it will be true for us to say that the judge's action is an act of mercy. And this will be true even if his action is also an instance of doing the just thing, i.e., giving the thief something she deserves. In general, whenever one is in a context in which the notion of permissibility that is relevant to whether x was merciful to y is *not* the notion associated with justice – i.e., the notion of what is permitted according to the constraints of justice – there is the possibility of one's truthfully saying that x was merciful to y . I suspect that this will typically be the case, for I suspect that the "default" notion of permissibility – the one that will be relevant in most contexts – is going to turn out to be the notion of moral permissibility.

Next, consider the puzzle about the wildly divergent intuitions about specific cases. According to FAM, there are two main explanations for this phenomenon. The first is simply that, according to FAM, it will sometimes turn out that people disagree over whether a particular action is a case of mercy because they are in different contexts, in which different notions of permissibility are relevant. For example, some lawyers discussing whether Case 1 (involving the judge and the lighter sentence) is a case of mercy will be in a context in which the relevant notion of permissibility is the legal notion; and so they will agree that Case 1 is indeed a case of mercy. But a group of moral saints, say, discussing the case will be in a context in which the relevant notion of permissibility is the moral notion, so that (assuming that giving the harsher sentence would have been morally wrong) they will agree that the case is not a case of mercy.

The second main explanation provided by FAM for the phenomenon of wildly divergent intuitions about specific putative cases of mercy is that in many cases, different people consider different notions of permissibility to be relevant to the question of whether x is merciful to y at t . E.g., Smart would probably consider the action of the victorious soldier in Case 2 not to be an act of mercy, whereas I have said that it is a case of mercy. The reason I am inclined to say that this is a case of mercy is that I am inclined to think that (in typical contexts in which we might discuss the case) the “wartime” notion of permissibility is the relevant notion. Meanwhile, if Smart insisted that Case 2 is not a case of mercy, as I suspect she would, then she would no doubt do so because she thinks that (in typical contexts in which we might discuss the case) the “justice” notion of permissibility is the relevant notion.

Similar remarks would apply to the disagreement between Rainbolt and myself regarding the case of the Nazi doctor. The reason I am inclined to say that case is not a case of mercy is that I am inclined to think that (in typical contexts in which we might discuss the case) the moral notion of permissibility is the relevant notion, whereas Rainbolt apparently thinks that the “normal” notion of permissibility is. Thus, a disagreement over which notion of permissibility is relevant to a given case can in turn lead to a disagreement over whether the case is one of mercy.

6 Objections

So far I have suggested that FAM, The Flexible Analysis of Mercy, is the most plausible theory, and also that it provides satisfying solutions to our two puzzles about mercy. I want to close by considering some objections to FAM.

The first objection involves cases in which an action that may strike us as merciful does not appear to be permissible under what is no doubt the relevant notion of permissibility.²⁶ One such case would be the following variation on Case 1 from above.

Case 6: A judge is sentencing a convicted thief. The law calls for anywhere from 5 to 10 years in prison, but the judge, moved by compassion for the thief, instead sentences her to just 3 years.

It is easy to imagine that the thief in this case would render a “merciful” verdict in Mercy Court, and in considering the case we are likely to say the same. It seems that the judge acted mercifully. But the judge’s action does not appear to count as merciful, according to FAM. The reason is that when discussing the case, whether in Mercy Court, ordinary life, or the seminar room, the notion of permissibility that will most likely seem relevant is legal permissibility, and the judge’s action in the case (sentencing the thief to 3 years in prison) is not in fact among the judge’s legal options. This means that the judge’s action will apparently not satisfy condition (iv) of FAM, which requires that the harsher alternative and the more lenient action actually performed *both* be permitted under the relevant notion of permissibility. Thus we seem to have a case in which an action that should count as merciful does not satisfy the conditions laid down by FAM.

I agree that this initially appears to be a counterexample to FAM. But I think that once we delve a bit deeper, we can see that it really is not. Here’s why. First, suppose we are in a context in which it’s clear that the relevant notion of permissibility is *legal permissibility*. For example, suppose we are members of a committee of lawyers charged with deciding whether the judge in Case 6 should be named Judge of the Year. Then the case is relevantly similar to Case 5 above (where Carla killed Alice, out of compassion for Betty, so that Betty would get the Nobel Prize). In each case, an agent does

²⁶ I am grateful to an anonymous referee for raising this objection.

something impermissible that constitutes less harsh treatment of someone else than an alternative that is available to the agent. But if I was right above to say that Carla's killing Alice should not count as an act of mercy toward Betty, since it was an impermissible thing to do, then the correct thing to say about this case (given that we are in a context in which the notion of permissibility that is relevant is legal permissibility) is that it too is not a case of mercy, for the same reason: acts of mercy must be permissible, and this act was not. As lawyers charged with deciding whether the judge deserves the Judge of the Year award, we should say that what the judge did was not okay, and instead of thinking about giving him an award, we should be thinking about bringing him up on disciplinary charges. Under those circumstances, it would be quite inappropriate to call what the judge did an act of mercy. And FAM seems to get the right result here.

Now suppose, on the other hand, that we are in a very different context, where the relevant notion of permissibility is *moral permissibility*. For example, suppose we are members of a committee of moral saints charged with answering the question of whether the judge himself is a moral saint. Then I think there are two importantly different scenarios to consider. Either the judge's action in sentencing the thief to 3 years was morally permissible (despite the fact that it was an illegal thing to do), or it was not. (Let's suppose for now that the harsher treatments of the thief that were available to the judge – sentencing her to anywhere from 5 to 10 years – were in fact morally permissible. (We will shortly consider a variation on the case in which this additional supposition is retracted.)) If the judge's illegal action of sentencing the thief to 3 years was morally permissible, then I think the right thing for a bunch of moral saints to say about the case is that the judge is in fact acting mercifully, just as FAM would suggest. But if on the other hand the judge's illegal sentence of 3 years was *not* morally permissible, then I think the right thing for a bunch of moral saints to say about the case is that the judge's action is not in fact an act of mercy (since it was not an okay thing to do), just as FAM would suggest. So either way, it seems to me, FAM gets the right result here.

One thing that emerges from our discussion of this case is the importance of considering, not just an example of some putatively merciful action, but a combination of an action paired with a context of evaluation. When an action is said to be merciful, details about the action matter a great deal, but the context of utterance can also make a difference. Another thing that reflection

on Case 6 reveals is how natural it is to take moral permissibility to be the relevant notion in most cases.²⁷ That's why, in initially considering the case, it was so natural to think that the judge's action, though illegal, was nevertheless an instance of mercy.

Here is a notable point about the connection between FAM and Mercy Court. In Mercy Court, the one agent is trying to convince the other agent that his or her action was a case of mercy. So according to FAM, one of the things that is at stake is whether the action performed was permitted under the relevant notion of permissibility. This raises an important question: *In the context of Mercy Court, what notion of permissibility is the relevant one?* I think that reflection on this question shows that facts about *the context of the action in question* play a huge role in determining what notion of permissibility is relevant in *the context of evaluation*. For example, suppose that we are dealing with a case in which a judge in New York City in 1975 was legally permitted to sentence a convict to anywhere between 5 and 10 year. And suppose she sentenced him to 10 years, the maximum sentence allowed by the law. Now let the judge and the convict be reunited in Mercy Court. If the best argument the judge can give is that sentencing the convict to life in prison was permitted under the Portugeuse penal code of 1675, then she is not going to get a favorable verdict in Mercy Court. And I think that what this shows is that in Mercy Court, which notion of permissibility is relevant will be largely determined by facts about the context of the act in question.

The second objection to FAM that I want to consider challenges the assumption – built into FAM – that the notion of permissibility relevant to *A* (the action that may or may not be an instance of mercy) and the notion of permissibility relevant to *B* (the harsher alternative that was available to the agent of *A*) must be the same notion.²⁸ Consider this example.

Case 7: Once again, a judge is sentencing a convict. The judge happens to know that the convict is an innocent man who has been unjustly framed. The law requires a 10-year sentence, but morality requires that the judge help the innocent convict to

²⁷ In fact, for the first ten years that I was working on this paper, I was convinced that MAM (The Moral Analysis of Mercy) was the correct account. It was only about eight years ago that I came to believe that FAM is in fact the right account.

²⁸ I am grateful to Elizabeth Harman for raising this objection in correspondence.

escape. So the judge, moved by compassion for the convict, does just that.

Suppose the case goes to Mercy Court. It is easy to imagine that the judge will win a favorable decision from the convict by arguing as follows.

Look, I could have very easily just sentenced you to 10 years, as the law required, and left it at that. But I knew you were innocent, and I felt bad for you, so I helped you to escape, at great personal risk to myself and my reputation.

What seems to be peculiar about this case is that there is no one notion of permissibility such that both what the judge did (help the convict to escape) and what she could have done (sentence him to 10 years) were both permissible under that notion. (For it was stipulated in the story that the law required a 10-year sentence, while morality required that the judge help the convict to escape.)

This seems to be a counterexample to FAM, since FAM's fourth condition requires that both the action performed as well as the harsher alternative that was available to the agent be permitted under the relevant notion of permissibility.

One possible response to this case would be to bite the bullet, and maintain that the judge's action in this case is not an instance of mercy. For the proponent of FAM could insist that if the context picks out *moral permissibility* then the action is not merciful because the harsher alternative was not permitted, and also that if the context picks out *legal permissibility* then the action is again not merciful because it itself was not permitted.

Another possible response to the case would be to accept that this is a counterexample to FAM, and to revise the theory accordingly. That is, someone who finds FAM initially appealing, but who is moved by this example, could adopt the following account of mercy.

The Doubly Flexible Analysis of Mercy (DFAM): *x* is merciful to *y* at *t* =df there are two alternative acts, *A* and *B*, available to *x* at *t* such that (i) *x* performs *A*, (ii) *x*'s performing *B* would have resulted in more harm to *y* than *x*'s performing *A*, (iii) *x* is motivated by compassion for *y* to choose *A* over *B*, and (iv) there are two notions of permissibility, *P1* and *P2*, such that (a) *P1* and *P2* are both somehow picked out by the context of

utterance, (b) *A* is permitted under *P1*, and (c) *B* is permitted under *P2*.

I would of course be delighted if any readers of this paper were to come away from it with the view that, in light of Case 7, DFAM is the correct analysis of mercy. But I personally like the first response to this example. Here's why. It seems to me that Case 7 nicely illustrates how easy it is to switch back and forth between two different notions of permissibility while discussing a putative case of mercy. Suppose we are (as before) moral saints who are trying to decide whether the judge in this case is also a moral saint. Then it will not seem important that the harsher alternative (the 10-year sentence) was legal. Instead, it will only seem important that that option was morally impermissible, and hence we will want to say that the judge's action is not an instance of mercy. Or suppose we are (as before) members of the bar who are trying to decide whether to give the judge an award for being Judge of the Year. Then it will seem like helping the convict to escape was just a mistake on her part, and not a laudable act of mercy. Thus what Case 7 actually shows, it seems to me, is that context really matters.

Here is one last objection.²⁹ I said above that according to FAM, Case 2, in which the victorious soldier releases her foe without killing him, will count as an act of mercy. I also suggested that this should be an uncontroversial result. But perhaps this is a controversial result after all. For it is not hard to imagine the vanquished soldier delivering a negative verdict in Mercy Court, especially if it might have been morally wrong for Victor to kill Vanquished. ("I appreciate your letting me go and all. But since it would have been morally wrong for you to kill me, I don't think that counted as an act of mercy.") This seems to show that Case 2 is not uncontroversial after all, and thus that FAM yields a highly questionable result.

My response is to admit that Case 2 is perhaps not as uncontroversial as I earlier suggested, but also to claim that this is a consequence of the fact that the original Case 2 was somewhat underdescribed. To see why I make this claim, consider a sharpened version of Case 2, in which we stipulate that killing the vanquished soldier would be morally wrong. And to sharpen things even further, suppose that we are again in a context in which we are moral saints debating whether Victor too is a moral saint. Then, it seems to me, the notion of permissibility that would be relevant to the question of

²⁹ Here too I am indebted to an anonymous referee for raising this objection.

whether Victor acted mercifully toward Vanquished would be *moral permissibility*, and someone who said, in this imagined context, that Victor's act was merciful would speak falsely (since the harsher alternative would in this case have been impermissible under the relevant notion of permissibility). This is of course exactly the verdict rendered by FAM concerning this version of the case, with respect to this context of evaluation.

Now suppose instead that we fill in the details of Case 2 in such a way that it is not obvious that Victor's killing Vanquished would have been morally wrong. And suppose we consider a context in which the notion of permissibility that will strike us as relevant is, say, wartime permissibility. (For example, suppose we are members of a war crimes tribunal, and are charged with evaluating Victor's behavior.) Then, I would suggest, it does all of a sudden seem clear that Victor's releasing Vanquished should count as an act of mercy (since the harsher alternative, killing her foe, was permitted under the relevant notion of permissibility). And this too is exactly the verdict rendered by FAM concerning this version of the case, with respect to this context of utterance.

I think what these different versions of Case 2 show is that it is in fact a crucial advantage of FAM over its less flexible rivals that FAM, unlike those rivals, can accommodate the context-sensitivity of different attributions of mercy.³⁰

³⁰ I'm grateful to Peter Forrest, Elizabeth Harman, Frances Howard-Snyder, Hud Hudson, George Rainbolt, Theodore Sider, Kadri Vihvelin, Edward Weirenga, and several anonymous referees for helpful comments on earlier versions of this paper. I'm also very much indebted to Ben Bradley, Fred Feldman, Owen McLeod, Eric Moore, David Waller, and Erik Weilenberg for discussions that originally got me interested in the topic of mercy.