

NOTE TO READERS – This is an early draft. Many things are missing, including some citations. My thanks to Tara Alberts, Kristen Bell, James Edwards, Richard Ekins, George Rainbolt, Hamish Stewart, Jacob Weinrib, Lael Weis, and audience members at the universities of Toronto and Oxford for their feedback. A much-revised version is forthcoming in *Philosophy & Public Affairs*.

ABSTRACT – A pardon is an act of mercy according to the law, but is a pardon mercy in an ordinary or genuine sense? What distinguishes a pardon from a lenient judicial sentence, which is not mercy by the law’s lights? These are questions about what mercy as it is understood in law has to do with mercy as it is understood outside of law, and about who in government acts mercifully and when, if indeed anyone in government ever does. Here I propose a general analysis of mercy, then bring that analysis to bear on government action. Three features of my analysis are noteworthy. First, almost all existing analyses say that mercy is unconstrained in a normative sense, but I argue that mercy is unconstrained in the way that arbitrary power is unconstrained. Second, although it’s often assumed that mercy must be motivated by compassion, I show that mercy only requires acting with the intention to benefit the recipient. Third, my analysis says that mercy requires the giver of mercy to overcome a motivation to treat the recipient harshly. Given this analysis, few government acts are merciful, but pardon is an institutional approximation or analog of mercy.

I. JONES’S CASE

In 1999 in Texas, Sharanda Jones was convicted of conspiracy to distribute cocaine. Under federal sentencing guidelines, Jones was sure to receive a long prison sentence. The sentencing judge had some discretion, though, and Jones had good reason to think that she would eventually be released. There was no physical evidence that she had ever possessed, bought, or sold cocaine. She had a young daughter, and co-owned a diner. This was her first offence, and it didn’t involve violence. But prosecutors pushed for “sentencing enhancements”: Jones was a leader of the conspiracy, they said, and she had perjured herself with “her false denials of guilt on the stand”¹. The judge was persuaded. He sentenced Jones to life in prison without the possibility of parole. She was 32 years old.

*Associate Professor, Faculty of Law, University of Oxford; Garrick Fellow and Tutor, Brasenose College.

¹ Wesley Bruer, “From a Life Sentence to Clemency From Obama”, *CNN*, September 2, 2016. Accessed July 29, 2017, <https://goo.gl/Jgxf7C>.

In 2015, when Jones was middle-aged, President Obama granted her a partial pardon, reducing her sentence to 16 years (the amount of time already served), and Jones was released soon after. On the same day, Obama pardoned dozens of other non-violent drug offenders. Over the next two years, he pardoned thousands more. The reason was simple: “The punishment didn’t fit the crime”.² Offenders like Jones deserved a punishment, Obama said, but not the severe punishments they had received. The law was too harsh, and it fell to him, as President, to correct for its excesses. Justice in a particular case notwithstanding a general rule to the contrary is ‘equity’. Obama acted equitably in Jones’s case. Over his time in office, he worked equity on a grand scale.

To pardon Jones, Obama used a power under the United States constitution known as the ‘clemency power’.³ Clemency is mercy in the public or political context, and merciful is how the Supreme Court of the United States as well as past Presidents have described exercises of the power. The ancestor of the clemency power is the ‘royal prerogative of mercy’, which is still a part of the constitutions of the United Kingdom and Commonwealth countries such as Australia and Canada. Pardoning powers in many other constitutions are often presented as powers to show mercy.⁴ Pardons are certainly merciful as far as the law is concerned.

Is the law’s understanding of mercy consistent with our ordinary understanding of mercy? The President showed mercy to Jones, in a legal sense, but was the pardon an act of mercy in an ordinary or genuine sense? That’s one question I want to answer. Here’s another: why would the law single out pardons as merciful? Is there anything importantly different about pardons compared with, for example, a lenient judicial sentence, which isn’t merciful by the law’s lights? I wonder, in other words, what mercy as it’s understood in law has to do with mercy as it’s understood outside of law. And I wonder who in government acts mercifully and when, if indeed anyone in government ever does.

To answer these questions, I’ll work from the general to the particular: I’ll propose an analysis of mercy, then bring that analysis to bear on government action. Three features of my analysis are noteworthy. First, almost all existing analyses say that mercy is unconstrained in a normative sense, but I’ll argue that mercy is unconstrained in the way that arbitrary power is unconstrained. Second, although it’s often assumed that mercy must be motivated by compassion, I’ll claim that mercy only requires acting

² “Commuting the Sentences of 46 Prisoners”, YouTube video, 2:13, posted by The Obama White House, July 15, 2015, <https://goo.gl/j9517N>.

³ United States Constitution, Article II, Section 2 (“The President ... shall have power to grant reprieves and pardons for offenses against the United States, except in cases of impeachment”). When this power is used to substitute a lesser punishment for a greater one, as in Jones’s case, it’s typically called a ‘commutation’. Sometimes ‘partial pardon’ is used instead. I’ll use ‘pardon’ for simplicity’s sake.

⁴ Constitution of Denmark, section 24 (“mercy”); Constitution of Indonesia, section 14 (“clemency”).

with the intention to benefit the recipient. Third, my analysis says that mercy requires the giver of mercy to overcome a motivation to treat the recipient harshly. Given this analysis, few governmental acts are merciful, though pardons closely approximate mercy.

I should say that my aim is a descriptive and conceptual analysis of mercy. There are important normative questions about mercy, including about whether mercy is always or sometimes permissible, obligatory, and so on. Similar questions could be asked about mercy in government. However, these normative questions presuppose an analysis of what mercy is and of when those in government act mercifully – and that’s the sort of analysis I want to provide.

II. STARTING POINTS

Both people and acts can be described as ‘merciful’. The two descriptions can diverge: a merciful person can fail to act mercifully sometimes, and even a person who isn’t merciful can act mercifully occasionally. Like other scholars, I regard mercifulness as exhibited by acts as more basic, so that’s what I’ll focus on.⁵ I also take it that, to act mercifully, you must act mercifully towards someone, where that person is the recipient of mercy. Thus, when we talk about merciful acts, the typical locution is something like ‘you are merciful to me now’.⁶

Standard examples of mercy in the literature involve crime and punishment, but mercy ranges widely. Here are three examples I’ll return to often. In the first, a creditor forgives a debt to spare an unfortunate debtor. Think of Shylock, had he listened to Portia. In the second, a pirate spares his captive, after the captive begs for his life. In the third, a victorious soldier spares his vanquished foe, in light of his desperate pleas. I do not say these are necessarily acts of mercy; it depends on how we fill the examples in. But these examples are all at least in the neighbourhood of mercy.

Philosophers disagree about what a merciful act is, but there’s agreement on three core conditions.⁷ For you to show me mercy now, first, you must have a *choice* between two acts, such that one act alleviates or prevents more harm to me than the other. I’ll refer to these acts as the ‘harsh’ and ‘lenient’ alternatives, for short. Thus, the creditor can forgive the debt or insist on repayment; the pirate can make his captive walk the plank or let him live; and the soldier can kill his foe or not. Second, you must show *leniency*: you must opt for the act which alleviates or prevents more harm. It wouldn’t be

⁵ See e.g. George Rainbolt, “Mercy: An Independent, Imperfect Virtue”, *American Philosophical Quarterly* 27 (1990): 169-173, at p. 170.

⁶ A point I take from Ned Markosian, “Two Puzzles About Mercy”, *The Philosophical Quarterly* 251 (2013): 269-292, at p. 271.

⁷ John Parrish and Alex Tuckness, *The Decline of Mercy in Public Life* (New York: Cambridge University Press, 2014), chap. 9.

mercy if the creditor insisted on repayment, the pirate killed his captive, or the soldier executed his foe.

It's easy to see that both of these conditions are satisfied in ordinary pardoning cases. In Jones's case, for example, Obama had a choice: grant or refuse the pardon. He chose to grant the pardon, thereby alleviating Jones's suffering. The same could be said of most pardons, so most pardons are at least potential instances of mercy. I say 'most' not 'all', because there are exceptions. Suppose I'm a former member of a mafia organization, turned state informer. While in prison, I'm safe. Were I to be released, I would immediately be killed by my former mafia associates. Granting me a pardon wouldn't be merciful, because it would harm me more than refusing me a pardon. Cases like this are obviously unusual, though, and they leave open the possibility that most pardons are merciful.

There's a consensus on a third condition of mercy: your choice whether to act leniently or harshly must have been, in a certain respect, *unconstrained* (or "not strain'd", as Portia said⁸). People express the same idea when they say that mercy is "discretionary"⁹, an "act of grace"¹⁰, or a "free gift"¹¹. You never 'have to' act mercifully. You always 'could' act harshly instead. Whether to act mercifully is therefore 'up to you'. Again, the condition is obvious in my examples: the creditor, pirate, and soldier are all, in some sense, unconstrained. There's also clearly a sense in which it is up to a head of state whether to grant a pardon.

When it comes to the unconstrained quality of mercy, however, the appearance of consensus is deceptive. That's because no one agrees about the sense in which mercy is unconstrained. What does it mean that mercy is an act of grace? What does it mean that it's a gift? In what sense did the creditor not have to waive the debt? In what sense could the pirate have made the captive walk the plank? In what respect was the vanquished soldier's fate up to his foe? These are all ways of asking about the sense in which mercy is unconstrained. This isn't the only difficult question about mercy, but it's the question which most divides philosophers, so it's a good place to begin.

III. UNCONSTRAINED BY JUSTICE

The best known answer is from Jeffrie Murphy. Murphy thinks that mercy is unconstrained by justice. Mercy is "never owed to anyone as a ... matter of desert or justice"¹². On the contrary, to treat me mercifully, you must have a right in justice to treat me harshly, which you waive in favor of treating

⁸ William Shakespeare, *The Merchant of Venice* (Delaware: Prestwick House, 2006), Act IV, Scene 1.

⁹ Parrish and Tuckness, *The Decline of Mercy in Public Life*, p. 252.

¹⁰ *United States v Wilson*, 2 U.S. (7 Pet.) 150 (1883), p. 160.

¹¹ Jeffrie Murphy, "Mercy and Legal Justice" *Social Philosophy and Policy* 4 (1986): 1-14, at p. 3.

¹² Murphy, "Mercy and Legal Justice", p. 3.

me leniently. When we add this suggestion to the first two conditions of mercy, we get this analysis:

(M1) Agent *A* is merciful to agent *B* at time *t* = df (i) at *t* there are two acts available to *A*, *x* and *y*, such that *x* would alleviate or prevent harm to *B* more than *y*; (ii) *A* has a right in retributive justice to *y* at *t*; (iii) *A* performs *x* at *t*.¹³

I've framed M1 in terms of retributive justice, because that's the sort of justice that Murphy has in mind when he talks about mercy in the sentencing context. So, an implication of M1 is that, if you can treat someone mercifully, then they deserved to be treated more harshly in virtue of their wrongdoing.

Murphy develops his analysis of mercy so that he can assess whether judges are able to act mercifully in their sentencing role. He says that judges have a duty to impose just sentences, not a right to do so. Justice is a judge's "job description"¹⁴. As a result, judges never act mercifully (in their official capacity). They don't act mercifully even when they treat someone equitably, by tailoring a punishment to fit the crime. "Judges ... who are unmindful of the importance of individuated response are not lacking in mercy; they lack a sense of justice"¹⁵.

What about heads of state, like Obama? Are they also under a duty to act justly? Murphy is cagey about this, but he allows that heads of state might not be constrained by justice. They "might legitimately draw upon values other than the requirements of justice and thus might legitimately ignore the just deserts of an individual"¹⁶. Peter Twambley, whose work Murphy relies on, is more definite. The state has a right (not a duty) to impose just punishments, Twambley says, and the head of state can waive that right on behalf of the state. Thus a President "has the power to grant a pardon, to exercise clemency ... in his (fictitious) office as right-holder"¹⁷. If Twambley is right, then Obama didn't act mercifully by pardoning Jones. Obama had two options: one harsh, to refuse a pardon, the other lenient, to grant a pardon. Granting the pardon would have been merciful only if Obama had a right in justice to refuse the pardon. But he didn't have that right, because Jones didn't deserve to stay in prison for the rest of her life. So, Obama didn't act mercifully, precisely because he acted justly.

M1 or something like it is probably the way that most philosophers think about mercy. But the analysis is open to at least three objections. The first

¹³ The basic schema I borrow from Markosian, "Two Puzzles About Mercy". I'm not attributing M1 to Murphy, because he would want to add that mercy must be motivated by compassion.

¹⁴ Murphy, "Mercy and Legal Justice", p. 10.

¹⁵ Murphy, "Mercy and Legal Justice", p. 8. Murphy takes a more nuanced view in Jeffrie Murphy and Jean Hampton, *Forgiveness and Mercy* (Cambridge: Cambridge University Press, 1988), pp. 177-180.

¹⁶ Murphy, "Mercy and Legal Justice", p. 9, fn. 17.

¹⁷ Peter Twambley, "Mercy and Forgiveness", *Analysis* 36 (1976): 84-90, at p. 87.

is pragmatic. M1 would make it impossible for us to know whether an act is merciful until we determine whether there is a right in justice to act harshly instead. However, it's controversial (to put it mildly) what people deserve in virtue of their wrongdoing, both in general and in particular cases. Were we to adopt M1, we couldn't say that you had acted mercifully unless we had already come to a view as to the normative status of your act. That gets things back to front. Much better to first develop a descriptive analysis of mercy, which we can then deploy to point out and distinguish the relevant features of acts. With those features identified, we can go on to assess the act's normative status. This is an incremental approach, where we work up to the normative issues that ultimately interest us. But it's possible only if our analysis of mercy is free of hard-to-apply normative conditions.¹⁸ (This objection isn't specific to M1; it's good against any analysis that casts the unconstrained quality of mercy in controversial normative terms.)

Even if it's right to understand mercy's unconstrained quality in normative terms, M1 goes astray in a second way. M1 says that you can't act mercifully unless you have a right in justice to act harshly instead, but counterexamples aren't hard to find. Think of the example of the victorious soldier, which I borrowed from Ned Markosian. Markosian says, and I agree, that the victorious soldier can act mercifully even if his vanquished foe didn't deserve to die.¹⁹ The pirate and his captive serves as a similar example: plainly justice doesn't entitle the pirate to kill the captive, but it can still be merciful for the pirate to spare his life.

A final worry starts from the fact that, were we to endorse M1, we'd be committed to saying that what is merciful is never just. Why would we be committed to that claim? Because it's implausible that treating someone leniently is just and that treating them harshly is just.²⁰ If a 16-year sentence for Jones is just, say, then life in prison is surely unjust. If you have a right in justice to treat someone harshly, then the lenient alternative is not just. By choosing leniency, you might act mercifully, but you won't act justly. This is a result that Murphy embraces. In his eyes, mercy and justice are foes, not friends.

However, it's common to think that an act can be both merciful and just. Consider the history of pardons. In England, the prerogative of mercy was traditionally used as an instrument of individualized justice. Through the 14th century, for example, most pardons in homicide cases were granted because of circumstances relevant to what we would now think of as the offender's culpability.²¹ Between the 16th and 18th centuries, the death

¹⁸ This paragraph is modeled on the discussion of normative versus descriptive analyses of lying in Thomas Carson, "The Definition of Lying", *Nous* 40 (2006): 284-306, at p. 288.

¹⁹ Markosian, "Two Puzzles About Mercy", p. 274, 276-77; see also George Rainbolt, "Mercy: In Defense of Caprice", *Nous* 31 (1997): 226-241, at pp. 228-229.

²⁰ Steven Sverdlik, "Justice and Mercy" 16 *Journal of Social Philosophy* (1985) 16: 36-47.

²¹ Naomi Hunard, *The King's Pardon for Homicide Before AD1307* (Oxford: Oxford University Press, 1969); Helen Lacey, *The Royal Pardon: Access to Mercy in Fourteenth-Century England* (York: York Medieval Press, 2009).

penalty was vastly expanded, eventually attaching to more than 300 offences. In this period, the prerogative was used for reasons we would find familiar: given the nature of the offence and the character of the offender, death was considered too severe a penalty.²² Even though pardons were used to fit punishments to crimes, and thus to do justice in particular cases, they were characterized as acts of grace and gifts of the sovereign – that is, as acts of mercy. We talk the same way today: Obama’s decisions to pardon drug offenders on equitable grounds were described as merciful by supporters and critics alike. When Jones’s lawyer heard that Obama had commuted the sentence, she said: “[Jones] has more than paid her debt to society and is truly deserving of the mercy she was given today”²³. M1 suggests that this long history of describing acts as merciful and just is conceptually confused. On the contrary, we should be suspicious of any analysis of mercy so at odds with ordinary usage.

These objections are enough to set M1 aside. But first let me say that I think Murphy got something very important right. He saw that a judge doesn’t act mercifully simply by taking into account all relevant factors and then imposing a lenient sentence.²⁴ Judicial equity, even in the direction of leniency, isn’t mercy. Suppose, for example, that a just sentence for Jones was 16 years, the time she had served when Obama granted her clemency. Suppose that the sentencing judge had discretion to sentence Jones to between 16 years and life in prison. And suppose that the sentencing judge had in fact sentenced Jones to 16 years, rather than to life in prison. By hypothesis, the judge would have acted justly. But he wouldn’t have acted mercifully. This is crucial, for my purposes, because it narrows the issue. It means that, if Obama’s pardons were merciful, it’s not merely because they were equitable and lenient. There’s something else going on. Whatever the missing factor is, it would have to be present when Obama commutes a sentence, but absent were a judge to initially impose that sentence.

IV. UNCONSTRAINED BY LAW

It will be tempting for legal scholars to try to improve on Murphy’s account by replacing justice with law. The idea is that you act mercifully towards someone only if, legally, you aren’t constrained to treat them leniently or harshly.²⁵ Using terminology more familiar in law, we could say that you act

²² JM Beattie, “The Royal Pardon and Criminal Procedure in Early Modern England”, *Historical Papers* 221 (1987): 9-22.

²³ Sari Horwitz, “President Obama Commutes Sentences of 95 Federal Drug Offenders”, *Washington Post*, December 18, 2015. Accessed July 29, 2017, <https://goo.gl/xQ7zuW>.

²⁴ A point made earlier by Alwynne Smart, “Mercy”, *Philosophy* 43 (1968) 345-359, at p. 349.

²⁵ See e.g. David Dolinko, “Some Naïve Thoughts About Justice and Mercy”, *Ohio State Journal of Criminal Law*, 4 (2007): 349-60. Dolinko is cautious about extending his analysis beyond the legal context (p. 360).

mercifully towards someone only if, legally, you have discretion to treat them harshly or leniently. That yields the following analysis:

(M2) Agent *A* is merciful to agent *B* at time *t* = df (i) at *t* there are two acts available to *A*, *x* and *y*, such that *x* would alleviate or prevent harm to *B* more than *y*; (ii) legally *A* has discretion to *x* or *y* at *t*; (iii) *A* performs *x* at *t*.

This analysis tells us that Obama acted mercifully by commuting Jones's sentence: he had the power to grant or refuse a pardon; the law allowed him to do either; and he acted leniently, by granting the pardon. Indeed, nearly any pardon will count as an act of mercy. As I say, this analysis will appeal to lawyers, because it's consistent with the law's own characterization of pardons as merciful.

There's an immediate problem with M2, however. In the last section, I imagined a judge who exercises his discretion to sentence Jones to 16 years rather than to life in prison. If the judge is permitted to impose either sentence, and chooses the lenient option, then M2 suggests he's acted mercifully. But, as I said, that isn't the case. The judge would have acted justly, but not mercifully. In general, judges don't act mercifully simply by doing their jobs.

But M2 can be changed to avoid the objection. Let's distinguish two orders of discretion. An act *x* is *first-order discretionary* if and only if it's permissible to *x* and permissible not to *x*. The act is *second-order discretionary* if and only if there are reasons for or against *x* which it is permissible to take into account and permissible not to take into account when deciding whether to *x*. The distinction I'm aiming at will be familiar to legal scholars. Suppose the legislature delegates to you the power to issue liquor licenses. You might be required to grant licenses only under certain conditions. You'll need to use your judgment to know when those conditions are met, but there's no discretion here. Or, you might be told to decide whether to grant a license based on a list of factors. You're not obligated to grant this or that license, so you have first-order discretion. Because you're not permitted to decide based on factors that aren't on the list, or to set aside ones that are, you don't have second-order discretion. The third possibility is that the legislature lets you decide which factors to take into account, and which to ignore or disregard. In this last scenario, you have first- and second-order discretion.

Sentencing judges typically have first-order discretion. The law doesn't tell them to impose a particular sentence. Instead, it sets out a range of sentences, from which the judge must choose. However, sentencing judges do not typically have second-order discretion. They are required to have reference to, for example, the nature of offence: did it involve violence? Was a weapon used? What was the harm caused? They must also have regard to the circumstances and character of the offender: was this a first offence? Has the offender shown remorse? What is the likelihood of re-offending? By contrast, the President has first- and second-order discretion. There's nothing he is required to take into account when deciding whether to grant

a pardon, and nothing he is required to disregard. The law lets the President set his own standards. (More on this below.)

Suppose we amended M2 to reflect the difference between the two orders of discretion. The analysis would say that you act mercifully only if you legally have discretion to x or y *and* discretion as to which reasons for and against x and y to take into account. That would mean Obama acted mercifully by commuting Jones's sentence. By contrast, the sentencing judge wouldn't have acted mercifully by imposing the same sentence on the same grounds. We'll have avoided the counterintuitive result.

Should we therefore accept a legal analysis of the unconstrained quality of mercy? I think that the distinction between first- and second-order discretion is important, and I'll return to it later. But even the revised version of M2 won't do, for three reasons. First, you can act mercifully absent legal regulation (in a state of nature, for example).²⁶ Second, we're often willing to say that someone has acted mercifully without knowing the legal status of their act. You don't have to know what international law says about killing enemy combatants to think that the victorious soldier can show mercy to his vanquished foe. Third, an act might be merciful even though the harsher alternative was legally prohibited. Recall the pirate. He prepares to make you walk the plank, but relents at the last minute. What he was going to do was unlawful, but it's still possible that he acted mercifully.

V. AT ANOTHER'S MERCY

So far, I've discussed attempts to show that mercy is unconstrained by justice and law. Justice and law are, of course, both normative standards. There are other normative standards, besides these two (e.g., fairness). In principle, an analysis of mercy could be constructed around any of them, such that an act is merciful only if it is permissible according to that standard. And, indeed, every existing analysis of the unconstrained quality of mercy takes this approach. If we could only identify the *right* normative standard, the thought goes, we'd understand what it means for mercy to be unconstrained. I won't go through the alternatives one by one. They're all vulnerable to the pragmatic objection I made against M1, for one thing. More importantly, the assumption is mistaken. For while it's true that mercy is unconstrained, that's not because it's unconstrained relative to some normative standard.

I take it as given that, to show mercy to someone, that person must be 'at your mercy'. Likewise, to receive mercy, you must be 'at another's mercy'. It would be natural to say that a debtor is at his creditor's mercy, that the captive is at the pirate's mercy, or that the victorious soldier has his vanquished foe at his mercy. And it would be jarring, I think, to say that you showed mercy to someone even though they were never at your mercy.

²⁶ Markosian, "Two Puzzles About Mercy", p. 279.

What does it mean to be at another's mercy? It is, roughly, to be subject to that person's whims, to be under their control. They can do whatever they like with you. They are in charge of your fate, not you. Think of prisoners and guards, wives and husbands in traditional societies, or children and parents today. In all these examples, one person is at another's mercy. As these examples suggest, being at another's mercy is often an unpleasant place to be. There is the potential for abuse and exploitation, subjugation and oppression. (I'll offer a fuller account shortly.)

Here's the point. For you to show me mercy, you must have power over me which is, in some sense, unconstrained. That your power is unconstrained, in whatever sense this is, establishes that I'm at your mercy. It's because your power is unconstrained that you can do what you like with me. But being at your mercy is not simply a matter of you being permitted to treat me well or poorly according to some normative standard. Being at another's mercy is something more brute, practical, and often ugly.

I can show this in two ways. First, imagine that someone is prohibited from treating you harshly according to every normative standard you can think of. It would be unjust, unlawful, immoral, unfair, imprudent, etc. for them to harm you. Surely you could still find yourself helpless before them, subject to their whims, and vulnerable to abuse and exploitation. Why? Maybe they don't know what justice, law, morality, etc. has to say. Maybe they don't care. The pirate is a good illustration. Normatively, his choice is clear: let his captive live. Even so, it's easy to imagine you, the captive, finding yourself subject to the pirate's whims, your life in his hands. And it's easy to imagine begging for your life – begging, that is, for mercy.

Second, and conversely, imagine that you're permitted to treat me harshly or leniently, according to whatever normative standard is thought to be relevant. Now imagine that I know embarrassing secrets about you. If you treat me harshly, I'll publish what I know about you. That's enough to keep you in line, and to ensure that you'll treat me leniently. My fate is in my hands. I'm not helpless, and I'm not at your mercy. Or, imagine that I'm a beguiler and a charmer. I can reliably manipulate you to do as I like. If you treat me harshly, it's because I let you think that's what you should do. I'm not at your mercy, whatever justice, law, fairness, etc. permits you to do.

In short: for you to show mercy to someone, you must have power to treat them harshly which is unconstrained, in some sense. The sense in which your power is unconstrained will show that the other person is at your mercy. But you can have me at your mercy while being constrained in a normative sense. And you can be unconstrained in a normative sense without having me at your mercy. For the purposes of mercy, therefore, the relevant sort of constraint isn't normative. Given that mercy *is* unconstrained, it must be unconstrained in some non-normative sense. In what sense, exactly? To answer that question, I need to explain in more detail what it means to have someone at your mercy. And to do that, I need to take a detour to republican thinking about freedom.

VI. ARBITRARY POWER

Republicans emphasize the importance of freedom as non-domination. To be dominated by someone is, in part, to be subject to their arbitrary will or power. What's crucial, for my purposes, is that republicans think of being subject to someone else's arbitrary power as equivalent to being at their mercy. Philip Pettit, for example, uses "exposure to the arbitrary will" of another and "living at the mercy of another" interchangeably.²⁷ This suggests an alternative way of thinking about mercy. It suggests that mercy is unconstrained in the way that power which is arbitrary is unconstrained. So republicans and I both have an interest in figuring out what arbitrary power is: republicans to better understand domination, and I to better understand mercy.

Republicans try to figure out when power is arbitrary by thinking about when power is non-arbitrary. Republicans agree that for power to be non-arbitrary it must be subject to practical "constraints"²⁸. These constraints must meet two conditions, best described by Frank Lovett.²⁹ First, the constraints must be *external* to the power-holder. They can't just be a matter of the power-holder's own judgment or psychology; they need to be backed by third-parties. Second, the constraints must be *effective*: they must ensure that the probability that the power will be exercised in accordance with the constraint is relatively high. Normative and legal standards don't count, unless they're backed by an enforcement mechanism. When we put these two points together, we get the idea that power is non-arbitrary only if someone else can force the power-holder to live up to a standard he or she didn't choose. Arbitrary power, on the other hand, is power subject only to the power-holder's own "will or pleasure"³⁰.

Some republicans, including Lovett, want to stop there. They think that the presence of external and effective constraints is enough to make power non-arbitrary. The content or substance of the constraints doesn't matter. Republicans who think this way are *proceduralists*. Others disagree. Pettit, most prominently, says that power is non-arbitrary only if the power-holder is forced to have regard to the interests of affected parties. That makes him a *substantivist*. I'm going to go with Lovett and adopt proceduralism here. It would take me too far afield to try to resolve this intra-republican dispute, for one thing. Also, I think that proceduralism better tracks what it is to be at another's mercy. If your power over me is externally and effectively constrained, then I'm not at your mercy, because I'm not subject to your

²⁷ Philip Pettit, *Republicanism: A Theory of Freedom and Government* (Oxford: Oxford University Press, 1997), pp. 31-32; Quentin Skinner, "Freedom as the Absence of Arbitrary Power" in *Republicanism and Political Theory*, ed. Cecil Laborde, John Maynor (Malden, MA: Blackwell, 2008), p. 94.

²⁸ Pettit, *Republicanism*, pp. 57-58; Frank Lovett, *A Republic of Law* (Cambridge: Cambridge University Press, 2016), p. 115.

²⁹ Frank Lovett, *Domination & Justice*, pp. 96-97; Lovett, *A Republic of Law*, pp. 115-116.

³⁰ Lovett, *A General Theory of Domination & Justice*, p. 96.

whims, and you can't treat me however you like. That's true even if you don't have to take into account my interests. Finally, I doubt that the choice between proceduralism and substantivism matters much for my purposes, and in footnote 43 I explain how to change my preferred analysis of mercy to accommodate substantivism.

Let me end with two clarifications. Republicans are interested in large-scale, social power. Understandably, when they think of constraints, they tend to think of general constraints, like rules and standards. But sometimes power is non-arbitrary because it's subject to more specific constraints. If I can blackmail you, then I can force you to do as I want. If I can beguile and charm you, then whatever you might independently want, you'll reliably do as I wish. In neither case am I at your mercy, as I said; and in neither case do you have arbitrary power over me, because my will constrains yours. The second clarification is that being subject to someone's arbitrary power is necessary but not sufficient to be dominated by that person. It's plausible that an additional condition is that the dominator has arbitrary power over some extensive range of activities of the dominated person.³¹ For that reason, you can be merciful to me without first dominating me.

VII. UNCONSTRAINED BY OTHERS

Back to the analysis of mercy. I said that mercy is unconstrained and that for you to show me mercy I first need to be at your mercy. Being at your mercy is not a matter of you being unconstrained by justice, law, morality, or whatever. It requires you to be unconstrained in a more practical sense. Now I can say what that means: I am at your mercy when you have power over me and there are no effective and external constraints on your use of the power.

My examples bear this claim out. The victorious soldier is free from meaningful outside constraints. It's wartime. He's on a battlefield. No one is in a position to force him to do anything. Whether the soldier kills his foe or lets him live turns solely on his preferences and judgments. Likewise, the pirate is free from effective social constraints. Indeed, the pirate is a paradigmatic example of someone living outside of society's rules. Whether the captive lives or dies is a matter of the pirate's pleasure. The creditor will be subject to various legal constraints (for example, he can only demand so much money in repayment), but not when it comes to whether to insist on repayment when a debt is due. That's a matter of his personal choice, which he's free to make in isolation. The law doesn't give anyone else a real say in that decision.

Here is how the analysis looks after incorporating my preferred account of the unconstrained quality of mercy:

³¹ Christopher McCammon, "Domination: A Rethinking", *Ethics* 125 (2015): 1028-1052.

(M3) Agent *A* is merciful to agent *B* at time *t* = df (i) at *t* there are two acts available to *A*, *x* and *y*, such that *x* would alleviate or prevent harm to *B* more than *y*; (ii) there are no external and effective constraints on *A*'s choice between *x* and *y* at *t*; and (iii) *A* performs *x* at *t*.

This analysis is the same as M2, except for (ii), which is the condition I've been working towards for the last three sections.

Earlier I distinguished the sort of discretion that sentencing judges and presidents possess. Although the law doesn't usually tell judges to impose a particular sentence, they must take into account all relevant factors, and set aside extraneous factors. If a judge fails to do so, it's a legal error. Likewise, if a judge passes a sentence for no reason at all, or just on a whim, that too is an error. These are not mere formal requirements. There is an elaborate system designed to ensure adherence to them, enforced by external bodies, including appellate courts. So judges don't exercise arbitrary power, and offenders aren't at a judge's mercy. It's no surprise, then, that offenders don't usually beg for mercy at the sentencing stage of a trial. The appropriate thing for them to do is argue for a lighter sentence, based on legally relevant factors. If they're denied, then they can challenge the judge's decision and try to claim the judge fell short of his or her obligations. Thus, according to M3, judicial sentences, however lenient, aren't merciful.

When it comes to pardons, the picture could hardly be more different. When presidents decide whether to grant a pardon, it's up to them what to consider, and what to ignore. They can grant a pardon for any reason they like, or for no reason. They don't have to explain or justify their decisions. And there is no judicial recourse if a president refuses a pardon. Thus the Sixth Circuit Court of Appeals held in *Sapp* that the "the very nature of clemency is that it is grounded solely in the will of the dispenser"³². He "may agonize over every petition; he may glance at one or all such petitions and toss them away."³³ Indeed, arbitrariness is a kind of hallmark of the clemency power. More than any other power in the American legal system, it can be exercised free from outside input or oversight.³⁴ So, aside from the exceptional cases I mentioned earlier, M3 is consistent with pardons generally being merciful.

What about other governmental acts, besides sentences and pardons? I can't think of any other state power which creates the opportunity to prevent or alleviate harm to specific individuals and which is also free of formal, legal constraints. So, if mercy is found elsewhere in government, it must be where formal constraints are ineffective in limiting the use of power. I suspect this

³² *In re Sapp*, 118 F. 3d 460, 465 (1997). The court is describing a governor's pardoning power, but the point holds true of the President, too. See, e.g., *United States v Klein*, 80 U.S. (13 Wall.) 128, 147 (1871).

³³ *Sapp*, 465.

³⁴ My claim is limited to pardons in the American legal system. Pardons in some other legal systems are subject to constraints imposed and enforced by legislatures and judges. See Andrew Novak, *Comparative Executive Clemency* (New York: Routledge, 2016), ch. 8.

is true of some low-level decision-making. Think of a health inspector. There are general policies and standards he is supposed to apply, but he can ‘find’ or ‘overlook’ infractions if he wishes. There’s no meaningful constraint on the inspector, at least for those without connections, resources, or savvy to hold him to account. In this scenario, some restaurant owners depend on the inspector’s good will. They have reason to humble themselves before the inspector to keep their businesses open. Such conditions are ripe for predatory behavior, corruption, abuse – and mercy. Something similar may be true of some police officers, prosecutors, immigration officials, and so on, though the real-life details will be complicated. Everything depends on whether, in practice, the constraints applicable to these actors make it highly likely they will act in accordance with them.

VIII. INTENTION TO BENEFIT

If I’m at your mercy, and you treat me leniently, do you necessarily show me mercy? M3 says yes, but that turns out not to be true. Imagine you and I are fighting a duel. You’re the better duelist, and before long, your sword is at my throat. I’m at your mercy. The blade begins to bite. Fortunately for me, at the last second you remember you left the iron on at home, and run off. All of the conditions of M3 are satisfied, yet you haven’t shown mercy. Why not? At this point, I need to go beyond the core conditions of mercy I set out earlier.

It’s common to assume that merciful acts are motivated by compassion or like emotions, directed towards the recipient. Compassion is certainly lacking in the dueling example. When you ran off, you were motivated by concern for your possessions, not by concern for me. So, if compassion for the would-be recipient is essential to mercy, that would explain why sparing my life wasn’t an act of mercy. However, as far as I know, no one has ever offered an argument for why compassion is essential to mercy. Compassion was often absent in ancient Greek and Roman thinking about mercy. Among contemporary philosophers, the claim is far from universal.³⁵ And, on reflection, I think we can see that compassion isn’t essential to mercy.

Consider two examples. In 1346, at the start of the Hundred Years War, King Edward III laid siege to Calais. The city held out for longer than anyone expected. When it finally surrendered, the King wanted to sack and pillage the city. Eventually, he was persuaded to accept a sacrifice instead: six burghers (leading citizens) would walk out of the city with nooses around their necks. And so they did. The mayor of the city led the way, and five volunteers joined him. The King, still angry after the long siege, ordered the burghers’ execution. But before the order could be carried out, his wife, Queen Philippa, intervened. She fell on her knees and begged: “I most

³⁵ See e.g. Smart, “Mercy”; Claudia Card, “On Mercy”, *The Philosophical Review*, 81 (1972): 182-207.

humbly ask as a gift, for the sake of the Son of the Blessed Mary, and for your love for me, that you will be merciful to these six men”³⁶. The King relented; the burghers were spared. This is a famous example of mercy, but the King didn’t act out of compassion. The King didn’t care about the burghers. He wanted to punish Calais, which had resisted him for so long. His wife convinced him otherwise, by appealing to what he did care about – her.

The example is not so unusual. People sometimes beg for mercy on behalf of others (daughters, husbands, friends, etc.). You can imagine the plea: ‘show mercy for my sake, if not for the sake of my daughter (husband, friend, ...)’. Even if you feel no compassion for the would-be recipient, these pleas might move you. If they moved you, you would have shown mercy, just as the King’s showed mercy to the burghers.

Someone might think that, while mercy needn’t be motivated by compassion for the recipient, it must be motivated by compassion for someone, like the Queen. Thus, compassion would still be essential for mercy. One difficulty is that compassion is a response to a person’s plight or suffering. Because it is the burghers who are suffering, not the Queen, compassion is not an appropriate emotion for the King to show towards the Queen. Here the objector might shift tactics: compassion is only one of several emotions which can motivate a merciful act, others being love and pity. The King acted out of love for the Queen, and that, my objector will say, is why he acted mercifully. However, I doubt that love for a third party, as opposed to love for the recipient, can motivate mercy. Think back to the dueling example. Suppose that, as you are about to kill me, you remember that you left your child in the back seat of your car. You dash off. You act out of love for your child, but this still isn’t mercy.

Consider a modern example next, in which compassion, love, and pity are all absent. In 2004 in Tehran, as Ameneh Bahrami was walking home from work, she was approached by a man named Majid Movahedi. Several times before, Movahedi had proposed marriage to Bahrami; each time she had refused. Now, as he approached, Movahedi was carrying a red container. He threw the contents in her face. It was sulphuric acid. Bahrami’s face was dramatically disfigured, and she lost both of her eyes. After the attack, Movahedi turned himself in to the police. The judge intended to sentence Movahedi to death, but Bahrami wanted *qesas* (retribution). “Inflict the same life that he inflicted on me”, she said.³⁷ The judge granted her request: Movahedi would be blinded. The sentence received international attention and condemnation.

In 2011, Movahedi was taken to Tehran’s judiciary hospital to have acid dripped into his eyes by Bahrami’s brother. (Bahrami, being blind, couldn’t

³⁶ Jean Froissart, *Chronicles*, vol. II, transl. Thomas Johnes (3rd edn., London: Longman, Hurt, Rees, and Orme, 1806), p. 226.

³⁷ Saeed Kamali Dehghan, “Acid Blinding Sentence Postponed by Iran after International Outcry”, *The Guardian*, May 14, 2011. Accessed July 29, 2017, <https://goo.gl/92vi5y>.

do this herself.) But at very nearly the last second, Bahrami lifted the sentence, as she was entitled to do under Iranian law. What was Bahrami's motivation? Her first reason was the international attention on the case. "The second reason I decided to pardon him", she said, "was because it seemed like the entire world was waiting to see what will happen."³⁸ Bahrami does not here (or elsewhere) give any sign of acting out of compassion or pity or love. But hers was an extraordinary act of mercy, hailed as such by observers in Iran and around the world.

What's going on in these examples? I think Andrew Brien has given the best answer so far. Compassion isn't essential to mercy. What's essential, Brien says, is acting with the right intention. Mercy "must be directed ... at alleviating the dire situation of the beneficiary", even if the "ultimate beneficiary is someone other than the person in the dire situation"³⁹. Brien's suggestion can account for all of the examples so far. In the Calais example, the King didn't care about the burghers. Nonetheless, he acted with the intention of benefiting the burghers, and a full explanation for his act would include the fact that letting them go benefited them more than executing them. Likewise, in the acid attack case, Bahrami didn't feel compassion for Movahedi. Even so, she acted with the intention of lessening the harm he would suffer. By contrast, when the victorious duelist runs off because she left the iron on at home, she benefits her opponent, but that's just a happy accident. It has nothing to do with why she acted as she did.

I need to make one clarification. Suppose you x with the intention of y ing. That means you intend, of your x ing, that it help bring about your y ing. But y ing need not be your final goal. You might intend, of your x ing, that it bring about your y ing, and in turn that it bring about your z ing. You can act with the intention of bringing about a further means *and* with the intention of realizing a goal. Thus, by saying that the King acts with the intention of benefiting the burghers, I don't deny that he acts with the intention of pleasing the Queen. The Queen feels compassion for the burghers, and wants to spare them; the King wants to please the Queen; so the King does what's necessary to spare the burghers. The King acts with the intention of sparing the burghers *and* with the intention of pleasing the Queen. Thus, it would be natural for a spectator to the drama to say something like: 'The King acted mercifully to please the Queen'.

Incorporating this new condition into M3 gives us the following:

(M4) Agent A is merciful to agent B at time $t = \text{df}$ (i) at t there are two acts available to A , x and y , such that x would alleviate or prevent harm to B more than y ; (ii) there are no external and effective constraints on

³⁸ Saeed Kamali Dehghan, "Iranian Woman Blinded by Acid Attack Pardons Assailant as he Faces Same Fate", *The Guardian*, July 31, 2017. Accessed July 29, 2017, <https://goo.gl/m7TvPk>.

³⁹ Andrew Brien, "Mercy, Utilitarianism and Retributivism", *Philosophia* 24 (1995): 493-521, at p. 501.

A's choice between *x* and *y*; (iii) *A* performs *x* at *t*; and (iv) *A* performs *x* with the intention of alleviating or preventing harm to *B* at *t*.

The addition of condition (iv) rules out some pardons as acts of mercy. Consider a revised version of the mafia informer example. As before, if I were released from prison, I would be killed. It's only in prison that I will live. Now, knowing of my past crimes, the President thinks I deserve to suffer more than he believes I'm now suffering. So, he grants me a pardon, hoping that I'll soon meet my end. However, unbeknownst to him, I find prison unbearable. Even death would come as a relief. By granting me a pardon, the President has benefited me. That wasn't his intention, though, so he hasn't acted mercifully.

XI. REVERSAL

I think M4 gives us most of what we want out of an analysis of mercy, but it's still not quite right, because it doesn't start far back enough. Suppose you have a son, who you love. As he grows up, things become almost unbearable. There are serious problems with drugs. The drugs lead to lying and stealing, and eventually to violence. After your son moves away, the problems get even worse. He takes your most cherished possessions and sells them, and this is only the first of many such betrayals. But your son still needs you. When he's in pain, he always turns to you, and you've never refused him. No one would blame you for turning your back on him, at this point. You won't. 'What kind of person would I be to forsake my own flesh and blood?', you think. And now, when he comes to you again, you help him, as he knew you would. There's something to admire here, and something to criticize, but this isn't mercy. When you do as you were always going to do and help your child, you show love, not mercy. Why is that?

Philosophical discussion is usually dry compared to the phenomenon being discussed. When it comes to mercy, the gap is a gulf, because mercy is exciting. Mercy marks a change of course, a disruption of the narrative. Think of how people talk about mercy. When you show mercy, your 'heart softens', and you 'relent'. Stock examples of mercy are full of people poised and ready to strike. There are pirates, bandits, and black knights. Guns are lowered. Hands are stayed. Lots of things happen 'at the last second'. Mercy often comes after someone pleads or begs for it, which shows the unconstrained quality of mercy, but also the urgency of the moment. Tales of mercy are suspenseful, too. Will the condemned prisoner receive a pardon? Will the vanquished foe be spared? What will happen to Antonio? Mercy is never a foregone conclusion. Mercy is dramatic. The story about the parent and the child, on the other hand, is meant to be tragic partly because the ending is never in doubt.

Few philosophers have tried to account for mercy's dramatic, disruptive quality. But if there's an exception, it's George Rainbolt. In "Mercy: An

Independent, Imperfect Virtue”, Rainbolt said that “a merciful act occurs when someone treats another less harshly than one would normally expect them to be treated under the circumstances”⁴⁰. For example, we don’t think that refraining from hitting people with your car is merciful. No wonder, Rainbolt said, because people don’t normally run over others when they drive. I’m not convinced. Suppose that most people *do* hit each other with their cars. You don’t, though, because you don’t want to ruin your bumper. That’s not mercy. Could Rainbolt reply that, in my example, it isn’t normal for *you* to hit people with your car, and that’s why you don’t act mercifully? Yes, but this highlights the indeterminacy of Rainbolt’s analysis. What’s normal depends on a baseline: what’s normal for you, what’s normal for you on days like today, what’s normal for you on days like today in the last year, etc. Rainbolt didn’t tell us how to set the baseline. Despite these criticisms, I think Rainbolt was onto something. I said that mercy is dramatic; but it’s more accurate to say, as Rainbolt did, that mercy is a departure from what’s expected.

Rainbolt eventually abandoned his first proposal. In “Mercy: In Defense of Caprice”, he put forward a second proposal. Roughly, the idea is that you are merciful to me if and only if you have a “fairly strong reason” to treat me harshly, but intentionally don’t.⁴¹ I’m speculating, but my guess is that Rainbolt thought this was a better way to capture the unexpected quality of mercy. You have a fairly strong reason to treat me harshly. It seems plausible that people tend to do as they have fairly strong reason to do. So, it’s likely you’ll treat me harshly. But you don’t, which upsets expectations. Or so, I imagine, his thinking went.

An initial worry is that Rainbolt’s second proposal seems to yield the wrong result in the parent and child example. You have fairly strong reasons to refuse help to your difficult, dishonest child. As a result, Rainbolt would seem to be committed to saying that you act mercifully by helping them. But that seems wrong, as I said. How could Rainbolt respond? When we speak of you ‘having’ a reason to act harshly, as Rainbolt does, we might mean that there is a reason for you to act harshly, whether you see it that way or not. But we could mean that there’s a consideration which you regard as a reason to act harshly, such that you’re motivated to do so. Suppose Rainbolt takes this second view. He could then say that you don’t act mercifully by helping your child because, while there’s a good reason to refuse help, you aren’t motivated to refuse.

Let me fill in this thought a bit. In the last section I said that, to act mercifully, you must intend to treat the recipient leniently. And if you intend to treat them leniently, then of course you’re motivated to do so. I’ve just said that, to act mercifully, you must be motivated to treat the person

⁴⁰ Rainbolt, “Mercy: An Independent, Imperfect Virtue”, p. 170.

⁴¹ Rainbolt, “Mercy: In Defense of Caprice”, p. 233. Rainbolt adds that it’s also enough that “most people in [your] circumstances would have a fairly strong reason” to act harshly, even if you don’t see it that way.

harshly. So, the picture is of conflicting motivations. You're motivated to treat the person harshly, but a competing motivation prevails, and you end up treating them leniently. According to this view, mercy occurs in the context of an internal tension or motivational conflict.

We're making progress, but we're not there yet, because competing motivations don't ensure disruption of an anticipated course of events. Suppose we're friends, and I break my leg. On Monday, you intend to go over to my house on Wednesday, to make me dinner. On Tuesday, you learn that I said some mean things about you behind your back. You're angry at me, but not so angry to abandon me. And so, on Wednesday, you do as you planned and make me dinner. Despite your competing motivations, there's no change of narrative, no drama – and no mercy. Now let's change the example slightly. Suppose that on Monday, before you make any plans to help me, you learn about the mean things I've said. You're angry. You don't want to help me. Then, on Tuesday, you think of me alone in my house with my broken leg and nothing to eat. You overcome your anger, and decide to help me. On Wednesday, that's just what you do. This is mercy.

The difference between the two versions of the example is the order in which the motivations arise. When the stronger motivation to act leniently precedes the motivation to act harshly, there's no mercy. When the stronger motivation to act leniently comes after the motivation to act harshly, there is mercy. The idea is that a merciful act is the culmination of a dynamic motivational process.⁴² Initially, you're motivated to treat someone harshly. Then, you check yourself, and as a result, act leniently. There's a shift in your 'dominant motivation', as I'll call it. That shift creates drama. It also explains why mercy is contrary to the expected course of events. Had you not controlled your anger, you would have treated me, the injured friend, harshly. In general mercy requires that you control or moderate an initially dominant motivation to treat someone harshly.

By 'motivation' I mean to include a range of mental states. One is the *desire* to treat someone harshly. A second is emotional states like *anger* or *hatred*, as in the injured friend example. But I suspect the most common initial motivational state in mercy cases is the *intention* to treat someone harshly. The victorious soldier had, until a moment ago, been trying and therefore intending to kill his foe. Sword raised high for the killing blow, he stops and turns away. It's mercy – and a reversal of intention. Shylock planned to take a pound of Antonio's flesh. Portia begs him not to. Had Shylock listened, he would have gone back on an earlier intention. Edward III ordered the burghers' execution; he reversed his order, at Philippa's request. Bahrami had demanded retribution; suddenly, she changed her mind. You beg the pirate for mercy. Why? Because he's not just a man on a ship, and that's not a diving board. He's a pirate, that's a plank, and you know what happens next. You beg because you're trying to change his mind. You

⁴² I thank Tara Alberts for conversations on this point.

beg to win a reprieve – and a reprieve is possible only if harsh treatment is planned.

The connection between mercy and the reversal of intention is brought out best in O. Henry's story, "A Retrieved Reformation".⁴³ Jimmy Valentine, an infamous bank robber, is released from prison. He travels from town to town, breaking into safes. Following him is Ben Price, a lawman. Things change for Valentine when he arrives in Elmore, Arkansas, where he meets Annabel Adams, the local bank manager's daughter. Valentine falls in love with Annabel. He proposes to her (she accepts), and he abandons his life of crime, changing his name to Spencer. But Price hasn't given up, and he finally tracks Valentine to Elmore. On the day he arrives to arrest Valentine, Annabel's nieces are playing in their grandfather's bank. One of them becomes trapped in the time-locked vault. Soon she'll die of suffocation. Even though he knows it will give away his identity, Valentine uses his burgling skills and tools to crack the safe, and the child is saved. Into this happy scene walks Price. At first it looks as if Valentine's new life will be ruined. This is how the story ends:

"Hello, Ben!" said Jimmy, still with his strange smile. "You're here at last, are you? Let's go. I don't care, now."

And then Ben Price acted rather strangely.

"I guess you're wrong about this, Mr. Spencer," he said. "I don't believe I know you, do I?"

And Ben Price turned and walked slowly down the street.

This is a story about transformation and mercy. It's also a story about a reversal of intention. Price plans to arrest Valentine. That's the whole reason why he's in Elmore. But when he sees who Valentine has become, he changes his mind, and walks away.

X. MY FINAL ANALYSIS

I'm going to use 'change of mind' to mean a shift in dominant motivation. That includes abandoning an intention, setting aside a policy, losing your feeling of anger, etc. With this piece of terminology in place, here is my final analysis of mercy:

(M5) Agent *A* is merciful to agent *B* at time *t* = df (i) at *t* there are two acts available to *A*, *x* and *y*, such that *x* would alleviate or prevent harm to *B* more than *y*; (ii) there are no external and effective constraints on *A*'s choice between *x* and *y* at *t*; (iii) at *t-2*, *A*'s dominant motivation is to *y*

⁴³ O. Henry, *Selected Stories*, ed. Guy Davenport (New York: Penguin, 1993). Also used by Carol Steiker as an example of mercy, but for a different purpose: "Tempering or Tampering? Mercy and the Administration of Criminal Justice" in *Forgiveness, Mercy, and Clemency*, ed. Austin Sarat and Nasser Hussain (Stanford: Stanford University Press, 2007), p. 16-17.

at t ; (iv) at $t-1$, A changes his or her mind; (v) at t , A performs x ; and (vi) A performs x with the intention of alleviating or preventing harm to B at t .

This analysis is the same as M4, except that it provides the ‘back story’ to the merciful act, in the form of conditions (iii) and (iv). The theme of M5 is the notion of controls or checks on power. When you can show mercy, it’s partly because you have power over someone which is unchecked by anyone else. You’re inclined to treat that person harshly. Then, you restrain yourself, for their benefit. You’re free from external checks, but exercise an internal check.⁴⁴

Why accept M5? Mercy is a concept in everyday use, and my analysis tracks our linguistic intuitions about mercy better than existing analyses. It correctly tells us that mercy may be found outside the criminal justice context. It explains how mercy is dramatic and disruptive. M5 also explains the relation between acting mercifully and having someone at your mercy, something which no other analysis does. My analysis suggests that mercy is morally serious, but it doesn’t beg any controversial moral questions. You can accept my analysis if you think that mercy is sometimes or always morally obligatory, prohibited, merely permissible, or supererogatory. You can accept it if you think that mercy mitigates a punishment which is just or unjust, deserved or undeserved.

Although M5 is very different than other contemporary analyses of mercy, it has much in common with how mercy was understood in the ancient world. John Parrish and Alex Tuckness explain how, within classical antiquity, mercy was understood “as treating an offender more leniently than one otherwise *would* have, rather than as one *should* have”⁴⁵. Seneca, for example, thought the essence of mercy is control of anger or the temptation towards of cruelty in the direction of leniency.⁴⁶ Thus, “an angry king who is about to punish with shocking cruelty and then restrains himself and punishes moderately” would be thought of “as showing mercy”⁴⁷. According to M5, this is exactly right way to think about the example – if a king, whose “anger overcomes all resistance”, “whom no man will interrupt”⁴⁸, nonetheless checks himself and shows leniency, then he is in that moment merciful.

M5 also helps explain what I think are our mixed feelings about the clemency power. The power is used to alleviate or prevent suffering, which is hard to object to. But the power’s existence means offenders are subject to

⁴⁴ If the substantivist republican account of arbitrary power is superior, then we should modify M5. In place of (ii), we would require that A is free of external and effective constraints which compel A to take into account B ’s interests.

⁴⁵ Parrish and Tuckness, *The Decline of Mercy in Public Life*, p. 54.

⁴⁶ Seneca, *Dialgoues and Essays*, transl. John Davie (3rd edn., Oxford: Oxford University Press, 2007), p. 214–215.

⁴⁷ Parrish and Tuckness, *The Decline of Mercy in Public Life*, p. 54–55.

⁴⁸ Seneca, *Dialgoues and Essays*, p. 193.

the state's whims. In republican terms, the clemency power fails Pettit's "eyeball test": those who need a pardon can't look those in power "in the eye without reason for fear or deference"⁴⁹. It also fails the "straight talk test": those who can grant a pardon have "reason to speak in the presumptuous tones of the master", and those who need a pardon "in the mealy-mouthed tones of the servant"⁵⁰. Offenders must humble themselves, beg for release, and pray for grace. This talk of deference and begging and grace smacks of the relationship between lords and serfs, not the relationship between a modern state and its citizens.

Finally, my analysis brings out an interesting connection between mercy and forgiveness. It's clear that mercy and forgiveness are different. You can forgive me without acting overtly, but mercy is about how you treat me. Forgiveness is usually thought of as a second-personal phenomenon (if I wrong you, then you can forgive me, but no one else can), but mercy can be third-personal. Still, mercy and forgiveness are often spoken of in the same breath, and it's worth wondering why. That brings me to another feature of forgiveness. Many philosophers think that for you to forgive me you must overcome some emotion, typically said to be resentment towards me.⁵¹ Mercy, I've said, involves something similar: a shift away from a negative attitude towards someone. The attitudes may be different (mercy doesn't require you to overcome *resentment* specifically), but in both cases there's an internal dynamic process.

XI. GOVERNMENT REVERSAL

What does M5 tell us about when those in government act mercifully? I'll start with pardons. The power to pardon is not subject to effective and external constraints, and in ordinary cases a pardon will prevent or alleviate harm to the recipient and be intended to do so. The issue is therefore whether, in ordinary cases, a pardon is a reversal of an intention or other motivation to treat the recipient harshly.

Sometimes a president will be motivated to treat some specific offender harshly. If he or she overcomes that motivation, and grants a pardon with the intention of benefiting the recipient, then it will be mercy, just as it's mercy if the 'angry king' overcomes his temptation towards cruelty. So, some pardons are merciful. But most of the time a President won't hold any attitude towards a specific offender. It's unlikely that Obama had any attitude towards Jones, say, before he decided to pardon her. As a result, a pardon will not usually follow a *personal reversal*.

⁴⁹ Philip Pettit, *Just Freedom* (New York: WW Norton & Co., 2014), p. xxvi.

⁵⁰ Pettit, *Just Freedom*, xxvii.

⁵¹ Paul Hughes and Brandon Warmke, "Forgiveness" in ed. Edward Zalta, *The Stanford Encyclopedia of Philosophy* (Summer 2017 edition), accessed July 29, 2017, <https://goo.gl/HJqo18>.

Consider, though, the fact that in ordinary pardon cases there *is* a decision which is reversed, namely, the decision to impose a relatively harsh sentence on an offender. A judge sentenced Jones to life in prison; Obama reversed that decision. Both the decision and the reversal – the sentence and the pardon – were acts performed by or within government. There was no reversal by an official of his or her own decision. But there was a *governmental reversal*. What should we make of that fact?

Here is one possibility. Government, some philosophers think, is an exercise in shared agency. Officials share a plan, the elements of which are laws and legal orders.⁵² The plan gives some officials the responsibility for developing sub-plans, which once adopted become plans of the group. This includes judges, when they pass a sentence. The shared plan also gives some officials the power to modify or reverse sub-plans, if they think it appropriate. This includes presidents, when they grant a pardon. On this way of thinking, pardons are a case of ‘collective reversal’. It would follow that most pardons are acts of ‘collective mercy’. (I’ve framed the argument in terms of shared agency, but we could reach a similar conclusion if government is a group agent, as some philosophers believe.⁵³)

I don’t dismiss out of hand the possibility that government is an exercise in shared (or group) agency. However, there’s no avoiding the fact that the notion of shared (or group) agency is controversial. The extension of shared (or group) agency beyond dyads and small groups is also controversial. The application to government a further complication. I can’t explore or defend these ideas here. Depending on the nature of agency in government, it’s possible that pardons are ordinarily merciful. But that’s not a conclusion we’re entitled to draw, given what we’re entitled to assume. Let me pursue a more conservative line of thought instead.

I propose we understand governmental reversal as an interaction between separate agents. One agent, the judge, decides to impose a sentence. A second agent, the President, modifies or negates that decision by granting a pardon. This isn’t a case of an element in a shared plan being reversed, or of an agent reversing its own decision. Hence, it’s not a case of mercy. Hence, pardons aren’t ordinarily merciful. But we can still say something interesting about the relationship between pardon and mercy. In the case of both pardon and mercy, there’s a disruption of an expected narrative. In both cases, there’s a dynamic process. The difference is that with pardons the process is distributed across multiple agents, rather than being internal to a single agent. Pardon is ‘deconstructed’ mercy: it’s what we get when we spread the elements of mercy across an institution. Pardon is the institutional analog or approximation of mercy. What the law calls mercy isn’t mercy as it’s ordinarily understood, but it’s no surprise that pardon and

⁵² Scott Shapiro presents the legal system as an exercise of shared agency in *Legality* (Cambridge, MA: Harvard University Press, 2011).

⁵³ See e.g. Christian List and Philip Pettit, *Group Agency* (Oxford: Oxford University Press, 2011), especially pp. 39-40.

mercy are thought of together. This is not as bold a conclusion as the one in the last paragraph, but it is much safer.

Earlier, I said that lenient judicial sentences aren't merciful. The reason I gave is that sentencing judges are subject to a battery of external and effective constraints. Now I can give a second reason: judicial sentences are initial decisions, not reversals of earlier decisions, whether by the judge or anyone else. Pardons aren't ordinarily merciful either, but lenient judicial sentences have fewer features of mercy than pardons do. That helps to explain why the law treats pardons as mercy, but refuses to see lenient judicial sentences the same way.