

**Legitimately Arresting the Innocent, and Other Puzzles about Officially
Inflicted Harm**

Daniel Viehoff (NYU, Philosophy)

*For NYU Law & Philosophy Colloquium. Early draft. Comments very welcome. Please do not circulate
or cite without permission.*

It is widely held that at least some officials have duties to follow the rules of properly constituted political and legal institutions even if, in doing so, they intentionally inflict harm on an innocent victim who ordinarily has a moral right against such harm. For instance, the following three claims about the criminal justice system seem to be accepted by many:

1. Innocent people normally have a right not to be intentionally harmed.
2. Our criminal justice system inevitably makes mistakes. For instance, police arrest innocent suspects, and courts condemn innocent people to severe punishment. These mistakes inflict intentional (and often very serious) harm on their victims.
3. Police officers, prison wardens, and other officials are (within certain limits) under a duty to do their part in enforcing the law even if they thereby intentionally harm people who are in fact innocent.

Yet these three claims, though not strictly speaking incompatible, are in tension. They aren't strictly incompatible: Perhaps the officer, though under an obligation to enforce the law, is nonetheless morally barred from acting on that obligation by the victim's countervailing right (and if he does act in line with his obligation, his violation of the victim's right is at best excused). Or the officer's institutionally imposed duty is in fact weighty enough to justify infringing the victim's right, as a matter of lesser evil. Or

perhaps people are (under the right circumstances) liable to suffer harm at the hands of officials even if they are innocent, and so no right is violated or infringed by the officer's action. But each of these claims is either very difficult to defend or incompatible with central features of our legal-political practices surrounding officially inflicted harm. The aim of this essay then is to clarify how the tension may nonetheless be resolved.

Resolving the tension is of threefold significance. First, and most obviously, whether it can be resolved is of immediate practical import to the officials who are expected to do what would normally violate another's right. By extension, it is crucial for understanding the moral situation of the victim, whose permission to defend herself against the official's act depends, according to common views about self-defense, on whether the official had an adequate justification for inflicting harm on her.

Second, discovering that the tension cannot be resolved would cast a significant shadow over our legal-political institutions and many policies they pursue. It is not uncommon for proponents of libertarianism to argue that governments should be much less involved in the lives of their citizens because such involvement is frequently wrongful. If government officials do things that ordinary citizens would not be permitted to do, and if we lack an adequate justification for the special rights that officials have to do these things, then it seems to follow that government officials are not morally permitted to do many of the things on which governmental action depends.¹

Finally, by identifying the (I will argue quite demanding) conditions under which our

¹ See, e.g., Michael Huemer's argument for libertarian anarchism, which combining two claims: (i) states lack general political authority, and a general right to coerce their subjects; and (ii) many actions of officials would be impermissible if undertaken by private citizens. But if the state lacks authority, then officials aren't fundamentally different from private citizens, and so they are not permitted to engage in many of the actions that we think of as distinctly within the state's purview: (Huemer 2013).

criminal justice system is legitimate, we acquire crucial tools for criticizing, and arguing for revisions to, the state institutions under which we live. In other words, the following discussion is not an apologetic endeavor that seeks to justify legal and political practices as they are. Rather, it is a critical project that starts from claims central to our existing practice and then shows how our institutions would have to be organized for these claims to hold true. Correlatively it suggests that, insofar as our institutions fail to be so organized, the moral rights and duties of officials (and those with whom they interact) may be quite different from what is widely assumed.

This essay focuses largely on the actions of officials in the criminal justice system. This is not to suggest that they are the only ones facing the problem addressed here. Health inspectors and tax collectors, for instance, may encounter structurally similar challenges. But I am inclined to think that the problem is especially salient in the context of criminal justice: the harms inflicted on the victims (and especially the innocent victims) are especially vivid there, and the victims' rights are not ones normally thought to depend to a significant extent on existing social institutions or conventions. (This marks a significant intuitive difference between the rights threatened by mistaken arrest and imprisonment on the one hand, those at stake in, e.g., tax collection on the other.) To this we must add that for many citizens the criminal justice system is the (ugly) face of the state that they confront most regularly² – much more regularly than they encounter health inspectors or tax collectors. So there is special significance to developing an account of the conditions that must be met to establish *the legitimacy of the criminal justice system*: Since the morality of punishment prohibits punishing the innocent, how can we justify a system that

² For an especially powerful account of the role that the criminal justice system plays in the lives of many black citizens in the United States, see (Alexander 2010).

in practice inevitably punishes (and otherwise leads to the intentional infliction of harm on) innocent people?³ How, we might ask, can such a system, and the duties and actions of the officials that constitute it, be justified to innocent people who end up being mistakenly arrested or punished?

1. Sharpening the Puzzle I: Privately Inflicted Harm

The puzzling character of officials' rights and duties to harm becomes clearest if we contrast their moral situation with that of a non-official engaged in similar activities. For instance, when a police officer arrests a suspect, S, the suspect is non-consensually restrained and her freedom restricted. Normally S has a moral right against being so restrained, and restraining her in this way requires special moral justification. Imagine an ordinary citizen, C, were to restrain innocent S without her consent. (Call this case *Private Restraint*.) Then S would normally be seriously wronged by C. She would also normally be permitted to defend herself against the wrongful restraint, even at the cost of imposing quite significant harm on C, including harm that is somewhat greater than that which C threatens to inflict on S. Finally, C would owe her an apology and compensation for the wrongful restraint. (Similar points could be made about the restraints associated with pre-trial detention or punitive incarceration, and other instances of harm intentionally inflicted by officials in the pursuit of their duties. To keep things relatively simple, I will in the following focus largely on the case of arrest.)

³ The question this essay asks is thus importantly distinct from the more commonly discussed question why the state is permitted to punish at all (and why perhaps only the state may punish crimes), which focuses on its punishment of guilty people. See, e.g., (Duff 2001, Bennett 2011, Tadros 2011). The analysis of professional ethical obligations in (Appelbaum 2000) comes closer to dealing with some of the issues discussed here, but does not dwell on the central question of the proper distribution of harm that is the focus of much of the following discussion. Perhaps closest in focus is Ronald Dworkin's discussion of the ethics of legal procedures, in (Dworkin 1985).

Or at least all of this is true (most of us agree) unless certain special circumstances apply.⁴ Matters would be different if S were liable to be restrained; if C had a lesser evil justification for restraining S; or if C's action were excused. Let me sketch each of these scenarios, so we can inquire whether any of them explain the special case of officially inflicted harm.

a. Liability

First, non-consensually restraining S would not wrong her if S were *liable* to be so restrained. Since we normally have a right against being so restrained, such liability is a matter of somehow losing the relevant right. (Notice that the content of the right that is lost may be narrowly understood: perhaps it no longer wrongs S if a *particular* person inflicts this harm on S for a *particular* purpose, though it would still wrong S if the harm were inflicted by another person, or for another purpose.) There is some disagreement about what can ground liability, and how far the notion of liability extends.⁵ But there is significant agreement on paradigmatic cases. For instance, if S has unjustifiably attacked (or otherwise wronged) another, and restraining (or, more generally, harming) her is necessary to appropriately respond to her wrong-doing, then S has lost the protection against non-consensual restraint (or harm) that all of us normally have.⁶ Consequently S would not be permitted to defend herself against such necessary and appropriate restraint, nor would she be entitled to compensation after the fact.

S's moral situation would plausibly be similar where, though S did not

⁴ Liability and lesser evil justifications are widely considered the standard ways of defeating rights against being harmed. For discussion, see (Rodin 2011). Being excused is the standard way in which someone may escape some of the moral consequences of violating another's right where that right has not been defeated.

⁵ For a discussion focusing on liability to defensive harm in particular, see (Quong 2012). For a broader conception of liability, see the discussion in (Tadros 2011).

⁶ This formulation is meant to leave open the precise relation between necessity and liability. For competing views, cf. (McMahan 2009) and (Tadros 2012) on the one hand, (Firth and Quong 2012) on the other.

unjustifiably attack another, S violated an obligation not to act in ways that give others good reason to believe that S poses a wrongful threat. Yet S can be asked to bear such an obligation only if it imposes a relatively small burden on her, and significantly reduces the risk of costly and faultless mistakes by others. Consequently such obligations are generally very limited: depending on context, very many actions can give others good reason for believing we pose a wrongful threat, and yet demanding that we abstain from all of these actions (or otherwise at least accept the risk of being defensively harmed) is asking too much of us.⁷ (To offer a concrete example: We can be expected not to walk around with a drawn gun, and if we do, we may have to bear the cost. But we cannot be expected not to walk around with a screwdriver. And yet someone carrying a screwdriver can, in certain contexts, reasonably be seen as posing a genuine threat.)

b. *Lesser Evil Justifications*

Matters are more complicated where, though S is not liable to be restrained, C has reasons to restrain S that nonetheless defeat S's right. (These are commonly thought of as *lesser evil* justifications.) For instance, if C's non-consensually restraining S is the only way to prevent D from killing an innocent child, then C may be justified in restraining S even though S remains protected by her right. Though C *infringes* S's right in this case, C

⁷ Another basis for making oneself liable for being restrained is to fail to fulfill one's duty of cooperation in solving a crime that one has not committed. S is morally obligated to explain her whereabouts, and perhaps empty her pockets, if she is in the vicinity of the crime scene and reasonably suspected of committing the crime. This is a relatively small burden she can reasonably be expected to bear to protect the interests of crime victims, even if she has no duty not to be in the place where she is found. If she fails to bear the burden voluntarily, she may be forced to bear it, and enforcing that duty may require arresting her. All of this is true independently of institutional background: even in the proverbial state of nature we have such duties to cooperate with our fellow denizens who need to protect themselves from criminals. But notice that in this case, the restraint should be released as soon as her pockets have been searched etc. The case this essay mainly focuses on is one that arises if something incriminating is found in her pockets, and she will thus be restrained for an extended period of time, at significantly greater cost to her, even though she is innocent; and the question it considers puzzling is how imposing this cost on her, despite her innocence, could be justified.

does not *violate* her right, since the infringement is justified all things considered.⁸ In such cases, S would normally not be permitted to defend herself against C's justified action. Nor would a third party aware of all the relevant facts be permitted to interfere and prevent C from restraining S. Yet S would usually be entitled to subsequent compensation, whether by C or (more likely) by D.

c. Excuses

But what if S is not liable to be restrained, nor is a lesser evil justification for restraining S available? Imagine, more concretely, that C reasonably but mistakenly believes that S is unjustifiably attacking another or her property (or has so attacked), and has thus become liable to be non-consensually restrained for defensive purposes.⁹ (I focus on this case on the assumption that something like the need to prevent wrongdoing against others is what underpins the laws that officials enforce when making arrests. I say more about this assumption in the next section.) Furthermore, C believes reasonably but mistakenly that restraining S is necessary and proportionate to protect another from wrongful harm, and so justified all things considered. In other words: Relative to the evidence available to C, it is reasonable to conclude that there is a threat-based liability justification for restraining S. But though this conclusion is warranted by the evidence available to C, it is false: S does not pose a wrongful threat to anyone (nor has S violated

⁸ For the distinction between infringement and violation, see (Thomson 1990)

⁹ Discussions of permissible harming often distinguish between 'eliminative harming' and 'manipulative harming', where the latter is usually taken to be more difficult to justify than the former. Arresting someone who is about to commit a wrongful act may look rather like eliminative harming, and so appear to be subject to a lower threshold for justified restraint. Arresting someone who has already committed a crime, or is suspected of having committed a crime, and whose guilt or innocence the police is investigating, is closer to manipulation. And punishment, if justified instrumentally, is most plausibly an instance of manipulative rather than eliminative harming. (For a discussion of this last point, see (Tadros 2011), Ch.12.) So at least parts of our practice of arrest seem to require justification at the higher threshold of manipulative harming, and it is on these that I focus here.

any obligation she may have to act in ways that make it reasonable to suspect her of posing such a threat).

In cases of this sort we may say that, though C's restraining S is not justified, it is nonetheless *excused*.¹⁰ The term itself is less significant than the fact that cases of this sort seem to differ in important ways from those where S is in fact liable to be restrained, or where restraining S is in fact the lesser evil. Where C's restraining S is merely excused, C wrongs S. In doing so, C acts non-culpably. And yet, I think, there is little doubt that S would usually be entitled to an apology and compensation. Furthermore, if S could prevent C's restraining her, she would, I believe, usually be permitted to do so if her resistance inflicted no more harm on C than the restraint would inflict on S – as would be a bystander who knows all the relevant facts.¹¹

2. Sharpening the Puzzle II: Officially Inflicted Harm

Return now to the case of police officer with which the previous Section started. Imagine officer P, acting either on his reasonable good faith judgment that S is committing, or has committed, a wrong (and has thus become liable to be restrained) or pursuant to a superior's order (itself based on a good faith judgment that S is likely committing, or has

¹⁰ There is a vast literature on the distinction between justification and excuse, both in morality and in the law. My concern here is not with the general distinction, but with the specific case at issue. What matters for the purposes of this argument is that officials may justifiably take a greater risk of error than private citizens (with all of the differences in liability to compensation and defensive harm this entails). This point stands even if it turns out that private citizens may justifiably take *some* risk of error in arresting people.

¹¹ There is some disagreement on this last point. Jeff McMahan, for instance, suggests that someone in roughly the position I attribute to C is an 'innocent threat', and that, though an innocent threat is liable to suffer some defensive harm, that liability may be too minimal to justify a serious inequality in the distribution of harm between C and S: (McMahan 2009), ch. 4. I doubt this is right at least in fairly standard cases where C merely defends himself. I discuss the complications that arise where C defends others below; for related concerns about McMahan's position, see (Tadros 2011), pp. 228-235.

likely committed, a wrong), arrests S.¹² S is, let me stipulate, innocent, and the judgment behind P's action mistaken. Nor has S engaged in action that she was obligated to avoid so as not to create reasonable suspicion. (She was not, for instance, obligated to avoid being in the vicinity of a crime scene while carrying incriminating tools, and yet this is the basis of P's arresting her.)

a. Liability

In core cases of arrest or punishment, the official acts on the basis of a judgment that the subject is liable to suffer the harm imposed. Arrest and punishment targets those who commit certain wrongs, and these are the kinds of wrongs that would, under many circumstances, make the wrongdoer liable to suffer certain kinds of harm for defensive or preventive purposes. (This is not to deny that there couldn't be cases of arrest and restraint that have a different basis. Protective custody may fall into this latter category. I treat these cases as exceptional, and set them aside in the following discussion.)

But sometimes the officials make a mistake: The person they are arresting or punishing has done nothing wrong. Consequently she is not liable to suffer arrest or punishment because she has done wrong. This does not as such entail that she is not liable to be arrested or punished. But if she is, it must be on another basis than that which seems to be what guides the official's action. (This is a familiar enough phenomenon where agents have a 'right to do wrong': the reason they take themselves to have to do x, and the reason why their doing x is permissible, come apart.)

But at least as liability is ordinarily understood, it is difficult to see what this other

¹² There may be considerations, having to do with the rule of law and the distinctive coercive powers of the police, that favor imposing particular constraints on permissible police intervention. I do not want to dwell on these issues here, and will simply assume that the wrong for which P arrests S is such that S becomes liable, not just for restraint in general, but for restraint by the police in particular.

basis could be. An agent's liability to harm, it is usually assumed, arises from an action of hers that is in some important way morally problematic. (Accounts of liability to suffer harm vary in filling in the details: the agent may be liable because she poses a threat to innocent others for which she bears culpability or responsibility, or which, relative to her evidence, was wrongful, or amounts to denying a certain moral status to others.¹³) But S has done nothing that poses harm to others; nor need she have had any reason to believe that her action poses such harm, or that others would reasonably take her action to pose such harm.

Nor is it clear that the law treats those who have been mistakenly arrested or punished as having been liable to suffer what they suffered. If the subject were liable to suffer arrest or punishment, she would lack the right that she would ordinarily have against being so harmed or restrained. But if she lacks this right, then she also lacks entitlements to compensation for infringements or violations of that right. And yet when the state discovers that, though everything was done by the book and a conviction was reasonable at the time it was handed down, the person convicted is in fact innocent of the crime, the state has a moral duty not just to release the wrongfully convicted but also to compensate her for (some of) the punishment she suffered. To be clear, not every time an official makes a mistake of this sort, the victim has a claim to compensation. (The difference between cases where compensation is due, and where it is not, will itself concern us later on.) But the fact that there are ever such cases indicates that the victim of official error

¹³ For a detailed discussion of the various options, see (Quong 2012). Let me flag here that Quong expressly limits his discussion to liability to defensive harm, and highlights that liability to suffer punishment, for instance, may be grounded in other considerations. But this doesn't help when it comes to making sense of, e.g., arrests, which are most plausibly justified on broadly defensive grounds. And it is an open question whether the best justification of punishment is not ultimately also based on what may be thought of as defensive considerations. For an argument to this effect, see (Tadros 2011).

cannot be generally assumed to be liable to suffer the costs of the official's reasonable mistake.

b. Excuses

Since the arrest flows from a reasonable but mistaken belief that S is not in fact innocent, it may seem natural to assume that P's arresting S is (just like C's arresting innocent S) an unjustified infringement of S's right, and P liable to pay compensation and suffer (at least some) defensive harm. In other words, we might expect our assessment of officials who arrest an innocent person based on reasonable suspicion to run parallel to how we view ordinary citizens whose arrest of an innocent person is merely excused (or what I will call *Excused Restraint*).

But that is not in fact how our moral, legal, and political practices treat the mistaken arrest of innocent suspects by the police. As Kent Greenawalt observes, "At least in the common law, it is universally said that police are justified [and not just excused] in making arrests based on probable cause... Rights of police officers to use force follow closely their rights to arrest, and people generally have no right to resist arrest or intervene to stop arrests based on the ultimate innocence of the person arrested."¹⁴ And officers who have reasonable suspicion¹⁵ that S has committed a crime are not normally expected to apologize, nor to provide compensation, for arresting innocent S. (And, for that matter, neither are police departments as institutions.) Even though innocent, S is expected to bear the cost of the officer's reasonable mistake and cannot pass it on to the

¹⁴ Cf. (Greenawalt 1998), p.23. I think the common law's treatment of these cases reflects, or at least fits, widely held convictions about the officers' moral situation.

¹⁵ US law distinguishes between reasonable suspicion, probable cause, and various other standards that have to be met to justify search, seizure, arrest, etc. Since my concern is not with doctrinal details, but with the philosophical foundations of our practice, I will speak generically of 'reasonable suspicion' even where the law would speak, e.g., of 'probable cause'. It should be clear that the standard of reasonableness may vary with context, including with the severity of the harm the official's intervention threatens to inflict.

officer.

Another way to think about the matter is suggested by the observation that those who are merely excused ought to regret doing what they did, and wish they hadn't done it; while those who are justified, though they may regret that the situation was such as to justify their act, nonetheless need (and indeed should) not regret doing what they did. Police officers arresting an innocent person based on reasonable suspicion fall, I think, squarely on the latter side of this divide. This is most obvious where P arrests S pursuant to a superior's order. P may regret that the order was issued. But given that it was issued, P need not regret making the arrest. Though perhaps less obvious, I think the same is true where P arrests S based on his own reasonable suspicion. P may regret that it was reasonable to suspect S despite her actual innocence. But he should not, I think, regret arresting her given that he had such reasonable suspicion (and no other admissible considerations against arresting her were in play).

c. *Lesser Evil Justifications*

All of this indicates that, although P's situation looks descriptively very similar to that of *Excused Restraint*, we do not think of officials like P as morally situated as if they were merely excused when they arrest an innocent person. We do, in other words, tend to think of officials like P as in a fundamentally different moral situation from that in which ordinary citizens like C¹⁶ normally find themselves when it comes to harming the

¹⁶ I leaving out one complication here: sometimes private citizens also have obligations to act on their own best judgment, even if it is mistaken, when it comes to harming another. For instance, doctors may be under a duty to perform surgery if they reasonably but mistakenly believe that the surgery is necessary. But it is important to recognize that these are exceptions within the domain of private action, reflecting the distinctive fiduciary obligations of doctors that they share with officials but not with most ordinary citizens. For a discussion of the role fiduciary obligations play in the common law governing harmful behavior, see (Thorburn 2008). (Tadros 2011), focusing on the ethics (rather than the law) of harm, also discusses the case of the doctor, without, however, highlighting the quite distinctive moral situation in which doctors find themselves qua fiduciaries.

innocent.¹⁷ But if we do not think of officials as merely excused, then it looks like the most plausible alternative is to think of them as justified in doing what they are doing. And once we set aside liability justifications, it seems the most plausible basis for justifying the officials' action is by appeal to lesser evil.

This is indeed how existing discussions of P's situation tend to try and solve our problem. They propose that some sufficiently important consideration must be available to outweigh S's right, and justify P's action. More specifically, central to some of the most influential discussions of the official's moral situation when ordered to harm an innocent victim has been the thought that the victim's right is outweighed by the official's countervailing duty of obedience. Thus David Estlund, in an important article, has argued that prison wardens, executioners, and soldiers have duties to abide by the outcome of a suitable decision-making procedure even if that outcome is mistaken and the victim innocent.¹⁸ As Estlund sees it, though an innocently convicted prisoner is indeed wronged by her imprisonment (and thus not liable to suffer it), the prison warden is "nevertheless morally obligated (and so morally permitted) to follow" the sentencing decision of a duly constituted court. The warden is in fact obligated to do so even if he in fact believes (correctly, and for good reasons) that the convict is innocent.¹⁹ And if the

¹⁷ Let me flag here another difference that I will not, however, discuss in detail in this paper: Where an official oversteps or abuses his authority, and thereby inflicts harm on another, the wrong committed would be more serious than might be similar harm inflicted by an ordinary citizen. For a discussion of officially inflicted harm that emphasizes this point, see (Gardner 2013).

¹⁸ (Estlund 2007).

¹⁹ Cf. Estlund's description of the moral situation of soldiers, to which he thinks that of prison guards is crucially analogous (p.215): "My thesis is precisely that under the right conditions, even though the victim is wronged by the unjustly warring side, the soldier on that side is nevertheless morally obligated (and so morally permitted) to follow all normally binding orders—those that would be binding at least if the war were just." (Estlund 2007) A similar position is endorsed in (Harel 2008): "A judge, a prison guard, or even an executioner is often entitled or obligated faithfully to execute the state's sentencing decisions. Such a duty ... is not boundless but it is much broader than the duties borne by a citizen. ... [A]n official ... is typically entitled or even required to perform this task [of punishment] irrespective of his private convictions concerning the appropriateness of the sanction..." (p.130)

warden is obligated to inflict such punishment, then police officers are presumably obligated to follow suitable orders that inflict much less harm on their innocent victim.

The appeal to sufficiently weighty duties of obedience provides a straightforward explanation for various features of the official's moral situation unaccounted for by the appeal to excuses, and for the thought that ordinary citizens would not normally have more than an excuse when they mistakenly restrain an innocent person. The official whose duty outweighs the innocent victim's countervailing right is justified in doing what he does even if he thereby infringes the victim's right. So he need not regret doing what he is doing. And because he is justified, the official may not be liable to suffer defensive harm at the hands of the victim. By contrast, an ordinary citizen, who is not under a similar duty of obedience, lacks an adequate lesser evil justification for similar actions.

One might worry that the appeal to obedience only explains part of the practice we are trying to make sense of: It may explain why P must arrest S if ordered to do so by a superior official, but it does not explain why P must (and, consequently, may) arrest S based on his own reasonable suspicion, nor why P's superior may order the arrest. But this problem is easy enough to fix, by adding that each of the officials has, among other things, been ordered to make or order arrests where they reasonably suspect that a crime has been committed.

Another problem faced by the lesser evil argument is more difficult to resolve: It is rather difficult to establish a sufficiently weighty duty of obedience, a duty powerful enough to defeat the rather strong rights we have against being non-consensually restrained or harmed. The rights we have against being non-consensually restrained are very strong indeed, as are the rights we have against being intentionally harmed. To justify their infringement requires significant harm on the other side – harm, crucially,

that is many times greater than the harm done to the victim who is protected by the right. Yet the innocent arrestee or convict normally poses no threat of harm to innocent others. Arresting or punishing an innocent person is not instrumentally beneficial in any obvious way that could figure in an acceptable justification of the harm inflicted. (The fact that punishing someone who is innocent will deter others is, for instance, widely taken to be an inadequate basis for justifying such punishment because, as it is often put, it ‘treats the victim as a mere means’. And though matters may be different where the instrumental benefit is enormous – many innocent lives will be saved etc. – it is difficult to see how this could generally be true of the actions undertaken by officials in the criminal justice system.) And non-instrumental arguments must satisfy the very demanding standard required to overcome the powerful protection provided by the innocent person’s right. So it is far from straightforward how the official’s duty could be sufficiently powerful to overcome the protection S’s right provides to her.²⁰

But rather than focus on the challenge of establishing a suitably weighty duty, I want to highlight how even such a duty (at least by itself) would fail to explain a central

²⁰ Consider Estlund’s own argument. On one reading, suggested by Estlund’s related discussion of authority in (Estlund 2008), an official who disobeys a suitably decision-procedure acts disrespectfully towards his fellow citizens, arrogating to himself the right to make certain decisions. But it is difficult to see how avoiding such disrespect could be sufficiently important to outweigh the harm done by, say, imprisoning for ten years an innocent person. On another reading, Estlund’s argument rests on instrumental considerations tied to the procedure’s epistemic reliability. (For such an interpretation of the argument, see (Renzo 2013)) But such an epistemic-instrumental argument is much more difficult to sustain, and the conditions on which its success depends are likely absent in at least a decent number of cases in which we nonetheless think officials ought to obey the order. (For a discussion of the conditions that have to be met for epistemic reliability to justify authority, see (Viehoff 2016).)

Finally, in both cases it is difficult to see why police officers should generally be shielded from defensive harm when they arrest an innocent person. Someone whose right is infringed as a lesser evil may not resist the infringement because such resistance would stand in the way of preventing the greater evil. See, e.g., (Tadros 2011). Yet P disrespects his fellow citizens by failing to try to arrest S, not by failing to succeed. But then S does not undermine the value of P’s showing suitable respect by resisting arrest. Similarly, if P is justified in arresting innocent S to minimize the risk of epistemic error, then S can further minimize such error by avoiding or resisting arrest if she is innocent. (Not everyone accepts that the victim is barred from defending herself in lesser evil cases. See, e.g., (Rodin 2011). But then we are obviously even further from justifying the practice at hand by appeal to considerations of lesser evil.)

feature of the practice at hand: it does not make sense of the compensatory practices associated with officially inflicted harm.

When the need to prevent greater evil justifies the infringement of another's right, the right-holder is normally nonetheless entitled to compensation for the infringement. You have a right that I not enter your property without your consent. Nonetheless I am justified in breaking into your holiday cottage if that is the only way to survive a severe snowstorm that catches me by surprise as I go hiking.²¹ And yet I must pay for the broken window, and perhaps for the use of your property, since you weren't liable to suffer this harm, and so your right was (admittedly justifiably) infringed.

By contrast, when P arrests innocent S pursuant to his superior's order, S is not normally entitled to compensation for the arrest. If the officer has reasonable suspicion (and otherwise does everything by the book), the infringement of S's right (assuming that is what it is) does not give rise to compensatory claims. This is not to deny that occasionally an innocent victim of legitimate official action is entitled to compensation. Thus in some jurisdictions a convict who has been exonerated may be compensated for her imprisonment even if there were good grounds for judging her guilty at the time of her conviction.²² Yet this is the exception, not the rule, where officials harm citizens pursuant to suitable directives or beliefs. Consequently a straightforward lesser evil justification of official obedience does not seem to fit most cases involving officially inflicted harm. A more fine-grained account – one that can explain why some, but not other, instances of such harm give rise to compensatory claims – is needed.

²¹ See (Feinberg 1978).

²² To be clear, even arrestees are sometimes entitled to compensation, when their arrest was not just mistaken insofar as they are innocent, but erroneous in the further sense that there was no reasonable suspicion in the first place. What is morally puzzling is why this is the exception rather than the rule.

The need for a more fine-grained account is also suggested by further reflection on the question of liability for defensive harm. I proposed earlier that lesser evil justifications do plausibly explain why officials inflicting harm on innocent victims pursuant to a suitably weighty duty are not liable to suffer defensive harm, and thus fits with how we treat criminal justice officials arresting or imprisoning innocent suspects or convicts. But notice that Estlund treats as analogous the cases of criminal justice officers on the one hand, soldiers on the other. And yet these are generally assumed to yield different forms of liability to defensive harm. Imagine that a police officer is directed to arrest S, who is reasonably suspected of murder. Even if S were innocent, S would not normally be permitted to defend herself against the arrest if doing so imposes harm on the officer. Contrast this with the case of a soldier who is sent into enemy territory to compel enemy forces to surrender. If the enemy's fight is just, their forces are clearly permitted to fight back against the soldier's attempt to restrain them.

I am not meaning to suggest that this observation by itself casts doubt on lesser evil justifications of the actions of criminal justice officials. But it does raise the question how we can account for the difference in liability to defensive harm that exists between such officials and soldiers. Someone might try to argue that the soldier's duty of obedience is not sufficiently weighty to defeat the enemy's right, while the police officer's duty is sufficient to defeat S's right. But why would that be? The enemy's right against being detained by our soldiers looks on its face no different from S's right against being arrested. Perhaps the thought is that the enemy soldiers are also at risk of being killed, while S is not. But we can stipulate that, if the enemy soldiers surrender without resistance, they face no risk of being killed. Conversely, if S were to resist the arrest with potential lethal violence, then the police officer too may be likely to use lethal force. What

about the duty that is meant to outweigh the right? It is, I think, difficult to see why the police officer's duty should be weightier than the soldier's duty of obedience.

If there is a difference between these cases, I want to suggest, it rests not on the *existence* or *strength* of the duty, but on its *ground*: on what justifies the duty's existence in the first place. The crucial difference between soldier and police officer is that the soldier's action against the innocent victim, if justified at all, is normally justified by the benefits it provides to the soldier's fellow citizens, but not to his enemy-victims. By contrast, the police purportedly protect all those within their jurisdiction, and thus act presumptively for the benefit even of those who end up as victims. As such we each have reason to bear some of the cost associated with policing, including the cost of being arrested subject to reasonable suspicion. Or so I argue in more detail next.

3. Mistakes and the Fair Distribution of Harm

Let me sketch in general terms the alternative account I want to offer, before developing parts of it in more detail in this and the next section. The central thought is that the official is acting on the basis of duties that are ultimately grounded in the benefits they provide to all the members of the community, S included. This in turn affects what burden S can fairly be expected to bear, and how S's claim to compensation for, or right to defend against, harm by officials may consequently be restricted. Indeed, it is only because S can fairly be expected to bear certain burdens as part of this scheme that the scheme itself can require officials like P to engage in actions that ordinary citizens would be free to avoid.

To understand the motivation for this view, it will help to return to the case of *Excused Restraint* that is descriptively (though not, we saw, normatively) closest to that of P's

arresting innocent S: the case where C is excused because he reasonably, albeit mistakenly, believes S to be engaged in wrongful action that makes her liable to be restraint for the sake of protecting others. If we can clarify why C's restraining S in this case would normally entitle S to take defensive action and claim compensation, we will be better positioned to understand which features of P's situation might explain the quite different rights S has where it is an official restraining her based on his reasonable yet mistaken belief. (Call this latter case *Official Arrest*.)

There is also a deeper methodological reason for returning to the case of private restraint. It is tempting to treat our intuitions about privately inflicted restraint or harm as morally basic. This in turn invites the thought that any deviations from the norms governing these cases are distinctly problematic, and especially difficult – perhaps impossible – to justify. For instance, many broadly Lockean views about political legitimacy start from intuitions about cases like *Private Restraint*, and reach the conclusion that political authority, and the rights of public officials to acts contrary to the rights citizens have against other private persons, cannot be justified except on the basis of the citizens' actual consent.²³ (Another way of putting this is that for the Lockean, the state, and public officials, are situated no differently than private persons, since private persons too could have authority and special rights to harm if the subject had given its valid consent.) And despite otherwise deep differences, the assumption that the norms governing *Private Restraint* will significantly shape what state officials may do also enters into many Kantian views of the state. They, however, take this as a reason for denying

²³ The most influential statement of this view can be found in the work of John Simmons: (Simmons 1979, Simmons 2001). For an explicit attack on the special rights of public officials that start from intuitions about what we (non-officials) may ordinarily do to each other, see (Huemer 2013).

that, outside of a system of law administered by public officials, we have full-blown rights that others not harm us where they have reasonable beliefs that we are posing a wrongful threat.²⁴

By offering an explanation of C's and S's moral situation by reference to more general moral principles and considerations, we may be able to dispel the temptation to treat cases like *Private Restraint* as morally basic, and any deviation from it as morally problematic, without yet having to deny outright the validity of common intuitions about the rights and duties that C and S have even outside of shared legal-political institutions.

a. *Fairly Distributing Burdens*

What explains that C normally incurs compensatory and defensive liability even though he acted reasonably in restraining S? I think the most plausible explanation for S's entitlement to defend herself, or claim compensation, is that she should not have to bear the burden of C's mistaken assessment of the threat she poses to him or others. Asking her to bear this burden would be unfair.

But why is it unfair? We cannot assume that S must never be asked to bear the burden of C's mistake. We saw earlier that at least sometimes (if the cost to S is low), we may bar S from engaging in activity that would reasonably foreseeably lead others to see her as a wrongful threat. In those cases, S would have to bear the cost of the others' mistake if she fails to abide by the prohibition. There is, I think, no good reason for believing that there couldn't also be cases where S is not barred from engaging in activity reasonably seen as threatening, but must bear the cost of another's mistake where that other does erroneously take S to be posing a threat. So we need an account that explains

²⁴ See, e.g., (Ripstein 2009, Stilz 2009).

why S need not generally bear the burden of C's mistake, even if asking her to do so is in principle possible (and, we will see later, indeed sometimes justified).

Thinking about cases where it would be fair to ask S to bear such a burden helps bring about a set of relevant considerations. First, it matters what kind of *choice* the parties had to suffer or avoid the harm at issue. If S is not required to engage in an activity, she had a choice not to do so, and may have to bear the cost of her choice where it leads to another's reasonably believing her to pose a threat. This is especially so if the other has relatively little choice in turn: If based on the evidence available to him, C reasonably cannot but conclude that S poses a threat, then C has no reasonable opportunity forming the belief that S poses a threat. It will then matter greatly whether C still has a reasonable opportunity to avoid acting on the belief.

A second crucial consideration is simply how much it will affect each person's life either to have to bear the burden, or to have to forego the choice at issue. If S has a choice, but the choice is of great significance for her capacity to lead an autonomous life, then asking her to forego the choice may be unfair. Conversely, if C lacks the choice, but that lack of choice is itself beneficial to him (say, because it is an upshot of a valuable relationship he has to another), then it might be unfair to ask S to bear the burden of C's lacking the choice.

In light of these observations, consider the costs and benefits associated with insulating an agent from defensive or compensatory liability for intentionally harming others based on reasonable yet mistaken beliefs. For the sake of simplicity I will say that in this case the agent has been *granted permission* to intentionally harm another based on reasonable beliefs. This may be misleading, since I do not mean to say that the agent is not obligated to avoid harming innocent persons even if she reasonably believes that they

are liable to suffer harm, and her action justified. (Nor do I mean to deny this. I wish to leave the issue open.) Nonetheless the practical implications commonly associated with such an obligation – liability to defensive harm and compensation – is absent. It is the absence of these liabilities that I mean to capture by using the shorthand of ‘permission’.

Being permitted to intentionally act to harm another based on a reasonable yet mistaken belief about another’s liability is normally beneficial to the agent who has the permission. Consider the case of defensive harm: Such a permission would allow its bearer to err on the side of protecting herself or others she cares about even if doing so imposes harm on someone who is innocent. By allowing false negatives, it minimizes the risk of false positives, to the benefit of the person who would suffer harm in the latter cases, and at a cost to the person who suffers in the former. Conversely, knowing that one’s action, together with another’s reasonable belief regarding that action over which one has no further control, will give others permission to inflict intentional harm on one will normally suffice to impose limits on one’s reasonable choice to take the action. Thus the fact that another has this permission is costly even where one ends up not being made to bear the actual cost of another’s mistaken belief. And where one does end up bearing these costs, the downside of being liable to suffer another’s acting on his mistaken beliefs is evident.

This explains why S could normally reasonably object if C were given permission to restrain her based on his reasonable but mistaken belief, and if she were concomitantly denied compensatory or defensive rights against C’s restraining her: C’s restraining her would impose an unfair burden on her, a burden that is unjustified by the benefit it provides to C or others.

Still, the argument from fairness requires further elaboration if it is to justify S’s

right to defend herself, or be compensated, in many cases where C reasonably but mistakenly deems her liable to suffer harm. Two points are especially important. First, we need a further account why some seemingly egalitarian distributions are nonetheless unacceptable. Second, we need to pay further attention to the fairness or unfairness of having to suffer harm at the hands of another who is acting on the basis of reasonable yet mistaken beliefs.

b. Equal Permissions to Act on Reasonable Yet Mistaken Beliefs

Tracing the rights and duties of the parties to questions about the fair distribution of benefits and burdens raises the question whether alternative moral schemes could not also be suitably fair. Thus someone might argue that, though it would normally be unfair to S if she had to suffer harm at the hands of C where C reasonably yet mistakenly takes S to pose a wrongful threat, such unfairness would disappear if C in turn equally had to suffer harm at her hands where S held a similar belief about C. So why not equalize the permission to defend oneself against wrongful threats based on reasonable beliefs alone?

The most obvious response might be that such competing permissions would lead to greater conflict, and that a scheme including them would thus be instrumentally harmful to all parties. But I am doubtful this is the best explanation we can offer. It is true that, if the permissions to act on reasonable yet mistaken beliefs are equalized, both parties are permitted to act on their reasonable yet mistaken perception of wrongful threats. C may thus be permitted to restrain or harm S (since C reasonably believes that S poses a wrongful threat to others), and S be permitted to restrain or harm C (since S reasonably believes that C is wrongfully attacking her). But whether this will in fact lead to an increase in the number of cases where both parties simultaneously believe that they are permitted to intentionally harm the other is far from obvious. It would depend on our

answer to questions such as these: How much more likely will agents be to act on their reasonable yet mistaken belief that another poses a wrongful threat if doing so is permissible, and does not give rise to compensatory and defensive liability, than if doing so is not permissible, and does give rise to compensatory and defensive liability, given that the agent already reasonably believes that the threat is actually wrongful and defensive action permissible? And anyway, might an increase in conflict here not be compensated for elsewhere? For instance, if acting on reasonable yet mistaken beliefs becomes permissible, then an agent doing so does not owe compensation for his mistake. Consequently the victim may now no longer reasonably believe that the agent owes her compensation, and so would not be justified in trying to extract such compensation from the agent. And thus one source of conflict – competing views about compensation that track competing views about the wrongfulness of the prior intervention – may have been taken off the table. In light of such considerations, I am doubtful that most of the work justifying the existing distribution of rights and permissions, and the rejection of equalizing permissions to act on reasonable yet mistaken beliefs, can be done by appeal to the instrumental costs of conflict.

A more plausible explanation is this: If two parties are both given permission to intentionally harm one another, then the result is that who ends up being harmed is determined by the morally arbitrary fact that one of them is physically stronger (or otherwise more capable of inflicting harm than) the other. Several reasons speak against a scheme that allows for this. One is that we generally do not think that our fundamental moral relations should be determined by arbitrary power advantages.²⁵ Another is that

²⁵ For a detailed defense of this intuition, see (Viehoff 2014).

the weak could plausibly complain that their bad luck in being physically weaker than others should play such a role in determining how they fare in life.²⁶ So we have general structural reasons to avoid moral principles that allow for the settlement of conflicting claims based on one person's greater power. Since equal permissions to intentionally harm others based on reasonable yet mistaken beliefs would have just this effect, we have reason to deny such permissions, and instead insist that an agent who acts on reasonable yet mistaken beliefs when intentionally harming another thereby incurs defensive and compensatory liability. To be clear, this is not to deny that as a matter of fact, agents whose beliefs about their respective liability conflict may end up fighting, and who prevails in that fight will be determined by that party's arbitrary power advantage. But it matters not just whether fights are so settled, but also whether our moral norms endorse or accept that they be so settled, and thereby imbue with normative significance the sheer fact that one side prevailed by force. It is this latter concern that explains why our moral norms are such that C's reasonable belief that S is posing a threat does not suffice to insulate C from S's defensive harm or claims to compensation.

c. Exceptions to the Rule

I have said that normally, an agent cannot expect others to bear the cost of her mistaken beliefs in their liability, and to suffer harm she intends to inflict on them as a result, without defending themselves against her, or claiming compensation. She cannot expect this because it would unfairly distribute benefits and burdens among the parties, most importantly by requiring her victim to bear costs that she herself had a fair opportunity to avoid creating. But this way of putting things in fact helps bring into focus

²⁶ This point is made by Victor Tadros: (Tadros 2011), p.208.

an exception to the rule: At least sometimes, someone lacks the choice to avoid imposing these costs on others. And where she lacks this choice, it may in turn become unfair to expect her to bear the cost in full.

Imagine a case where C reasonably but mistakenly believes that S poses a threat to an innocent third party, B. Ordinarily C is under no obligation to intervene in this case: our duties of rescue are relatively limited, and given the risk that intervention normally carries, C would usually be permitted to walk away. This reflects the significance of letting each of us pursue, within certain limits, her own projects rather than the projects of others. But if C's intervention as against S is not mandatory, then we might reasonably expect C to bear (much of) the costs where his exercise of discretion misfires. Otherwise S would be roped into C's discretionary project to protect others, and that itself would be a threat to S's independence. So if, despite having discretion, C decides to intervene and harm S, then S could reasonably insist that she should not have to bear the cost of C's choice to exercise her discretion one way rather than another.

But now imagine instead that C is under an obligation to intervene where B is wrongfully threatened. If the evidence available to C rationally requires her to believe that S poses a wrongful threat to B, then C lacks a moral choice in the matter. And if C lacks a moral choice, then it may seem unfair to expect him to bear the full cost of intervening. The harm must fall somewhere, either on S or on C; and why should C, rather than S, bear that cost if neither C nor S could reasonably have chosen to act in a way that helps avoid the creation of the harm?²⁷

Pointing to the mandatory nature of C's action is not, however, sufficient to

²⁷ This point is also made by Victor Tadros: (Tadros 2011).

explain why S should have to bear part (let alone all) of the burden. It also matters whose interests ground C's duty to protect B. Consider, for instance, a case where C has entered into a promissory obligation to defend another person, B, against wrongful attacks by any third party.²⁸ (Think of C as B's bodyguard or 'private protection agency'.) C reasonably but mistakenly believes that B is wrongfully attacked by S, and so C takes himself (reasonably but mistakenly) to be required to restrain S. There is, I think, little doubt that S's right to compensation and self-defense are the same here as in a case where C's action was discretionary.

Does this simply reflect the fact that C's duty is itself conditional on his discretionary exercise of his promissory power? Though I think it plays some role, I doubt it fully explains this case. Imagine C's duty to protect B is grounded, not in C's promise, but instead in a long-standing friendship that ties C to B. Imagine too that B and C grew up together, and have long history, etc.; and that at no point did C consciously choose to become B's friend. C's discretion is thus plausibly thought to play a limited role in creating the friendship. Yet he is B's friend, and may therefore owe B certain duties, including (perhaps) duties of protection against attack by others. Here too, I think, S could not be expected to bear the burdens that flow from B and C's friendship, or at least not where that friendship imposes such significant cost on S. Asking S to limit her self-defense against C if C takes himself to be required to attack her due to his friendship to B is expecting too much of her – not because C had discretion, but because the duties that C has are ultimately part of a relationship that is to C's (and B's) benefit.

²⁸ I think what I say next would also apply where C promises to protect B against any threat C reasonably perceives to be real and wrongful. Since I am not sure that such a promise would be binding in the first place, I focus on the simpler case at issue here.

These observations suggest a general strategy for explaining when S may have to bear burdens imposed by C's restraining or harming her even if C's action is based on C's mistaken belief: Whether S may have to bear burdens associated with C's acting on a duty to do what harms S depends on whose interests ground C's duty, and whether S can be expected to bear burdens associated with those interests. Two cases are of special significance.²⁹

i. Duties to Protect Third Parties: Lesser Evil and Liability

Our duties to rescue or defend are, I said earlier, limited. But where the numbers are large enough, we may have such duties even if they come at a significant cost to us. Imagine, for instance, that 10,000 innocent lives are at stake. In that case, I would almost certainly have a duty to take steps to rescue them, including steps to harm – even lethally harm – someone who poses a threat to these 10,000 lives. Now imagine that C reasonably but mistakenly believe that S poses a wrongful threat to 10,000 innocent lives, and that the only way to save them is to kill S. And imagine too that the only way for innocent S to avoid the threat posed by C is to kill C. May S defend herself against C's attack by killing him? Whether she may do so depends on what a fair distribution of the burdens at issue would look like. By assumption, neither C nor S had a reasonable opportunity to avoid

²⁹ I am setting aside a third one, which appeals to special obligations. If C harms innocent S to protect B, then it cannot matter whether C does so because he has a duty to protect B – that was the upshot of the earlier discussion of the bodyguard. What if both C and S have special obligations to protect B, and thus special duties to bear a burden for him? Even then, as long as C's obligation to protect B rests on a special relationship that in part benefits C, then S need not accept bearing a special burden where C makes a mistake. Imagine that C and S are both friends of B's. C mistakenly believes that S is attacking B, and steps in to restrain S. Since the question here is how to distribute the burden that flows from C's mistake between C and S, and S has no reason to bear a special burden for C's sake, she can wholly avoid the harm that needs to be distributed between them. Finally, imagine C and S have a special relationship. Here I do believe that, where C has a duty to protect B and mistakenly takes S to be attacking B, S may have to bear some of the burden of C's mistake. But this simply reflects the fact that we must generally bear some burden for our friends. This is, I think, largely independent of how reasonable C's action is, or indeed whether C had a duty to protect B.

the situation that they find themselves in, a situation in which one or the other will be killed. C couldn't avoid it because, based on the evidence available to him, he reasonably took himself to have no moral choice but to intervene and attack S. And S couldn't avoid it because she did nothing to bring about the situation in the first place. At the same time, the duty that made it unavoidable for C to intervene and impose the threat of harm on S rests on interests that make demands not just on C, but equally on S. (S has the same duty C has to intervene for the sake of protecting 10,000 innocent lives, and can, like S, be expected to bear a burden for the sake of the large number of innocents.) So as between S and C, each has an equal claim to avoid the harm that must now be inevitably inflicted on one or the other, and each is equally morally exposed to having the harm inflicted on him or her. I leave it open how best to respond to these equal claims in practice. Perhaps the best solution would be to require S to flip a coin, and determine on that basis whether to defend herself by killing C (and live) or not defend herself (and die while letting C live).

C's duty to defend the innocent thus exposes him to significant risk. Imagine C has a duty to defend some group G, based on the number of innocents in that group. C cannot reasonably avoid making a mistake in determining whether someone, S, poses a wrongful threat to G. In this case, C may have no adequate option to avoid inflicting harm on S up to the level that is necessary to effectively protect G. Neither S nor C had a reasonable opportunity to avoid the creation of harm that must befall someone among them. Often there may then be no basis on which to determine that one rather than the other can be fairly asked to suffer the harm.³⁰ But then, whatever level of harm C would

³⁰ Sometimes there may be further considerations that help determine who must bear the burden: Imagine S *knows* that C mistakenly but reasonably believes she poses a threat to G. She may thus also know that C reasonably but mistakenly takes himself to have permission to inflict up to some very high level of harm on S – say, lethal harm. But as it stands, less than lethal harm (say, restraining S) would be sufficient to fulfill

be imposing on S, S may be permitted to at least take steps to share the risk of such harm with C. And if S is permitted to take such steps, then C is exposed to (some share of) the risk that he is in turn imposing on S. And he is exposed to this share only because he had a duty to rescue or defend. Consequently whether he can be held to have that duty to rescue or defend depends in part on whether it is fair to expose him to a share of the risk that he is imposing on S as a result of having the duty. Because C has done nothing to incur liability to suffer harm for the sake of the innocent others, the level of risk he can be exposed for is significantly lower than that to which we can expose someone who has incurred such liability. But then it becomes difficult to see how C could have a duty to take defensive steps against S *as if S were* liable to suffer defensive harm. For then C would also run the risk of being exposed to harm at the level which he may exposes on a liable person, and yet he himself is not liable.

Public officials, like police officers, normally harm and restrain people on the basis of judgments about their liability to be harmed or restrained, rather than because they think such harm or restraint is a lesser evil, necessary to prevent much greater harm to others. But, as we saw, someone who mistakenly arrests or restrains someone by appeal to the lower standard of liability may have to bear a significant risk of harm at the hands of the innocent arrestee. Yet the public officials themselves have done nothing to incur liability to suffer such harm (and indeed are usually assumed not to be liable to suffer defensive harm even at the hands of those they mistakenly restrain or harm). But then it is difficult to see how we could justify the duties (and rights) of public officials by simply appealing to the large number of innocent people whom the official is trying to protect

the duty, as long as S does not resist or defend herself. In such a case, S would be required to bear the burden of restraint, since she is the only one positioned to prevent escalation.

from harm.

ii. Harm and Service

But C's duty to act in ways that harm S need not always solely rest on the interests of third parties. Sometimes C's duty may also rest, at least in part, on the interests of the person who ends up being harmed by C's acting as the duty requires. Imagine, for instance, that C is an old friend and neighbor of S's family. S has moved away, and S and C haven't seen each other in a long time. When S's parents die, the family home is empty for some time. In S's absence, C takes it to be his neighborly duty to protect the house from potential intruders. So when he notices a broken basement window, and sees a woman he doesn't recognize inside the house, C enters and restrains the apparent intruder, who is in fact S.

Ordinarily S could inflict significant defensive harm on another who restrains her inside her own house. She could do so even if that person reasonably but mistakenly believes that S is liable to be restrained. Imagine the person restraining S reasonably but mistakenly believes that this is in fact his house (the houses look very similar, and some prankster changed the house number), and that S is trespassing. I believe S could reasonably defend herself against such an attempt to restrain her, and inflict defensive harm on the would-be restrainer. This is true even if the person who mistakenly believes that this is his house also (rightly) believes that he has a duty, derived from the need to protect his children, to restrain any intruder in his home.

The matter is, however, different where S suffers such restraint at the hands of someone whose duty to restrain intruders in this house is itself grounded in a concern for S's interests. If C already incurs significant costs and risks of harm to protect B's home against intruders, then it would often be unfair to also impose on him the further cost of

any reasonable mistake he makes when acting on his duty of protection. This is not just to say that S and C would in this case be situated symmetrically, both equally morally exposed to suffer the risk and equally entitled to escape it. (This is, remember, the situation in which they find themselves where C's duty rests on the importance of protecting large numbers.) I think here, S may in fact have to bear the burden of C's mistake even if both S and C acted reasonably, and neither had a reasonable opportunity to avoid the scenario they now find themselves in. Though they are symmetrically situated when it comes to the opportunity to reasonably avoid the harm, the fact that C's duty ultimately rests on S's interests makes it fair to ask S to bear burdens that we do not normally think the victims of intentional harm based on reasonable yet mistaken belief can be asked to suffer.

4. Harm, Service, and Institutions

The discussion in the previous section provides us with the resources to explain the special rights and duties that public officials have to restrain or harm innocent people based on reasonable yet mistaken beliefs.

The starting point of this explanation is the observation that public officials have obligations to act for the sake of others that ordinary citizens lack. They are required to defend or rescue others even at a significant cost or risk to themselves. And, crucially, their obligation of defense or rescue extends equally to all members of the community, independently of the identity of those in need of protection. As John Gardner has put the point with regard to police officers: "Police officers, as officers of the law, are there to protect the bad guys as much as the good guys, except to the extent that specific protections are withdrawn by the legal mechanisms of accountability to the law itself

(arrest, search, detention, charge, remand, trial, sentence) or by the legal recognition of extreme exigencies with which the ordinary legal mechanisms (of arrest, etc) cannot be expected to cope. Subject to these strict exceptions, the police are there to protect looters against shopkeepers no less than shopkeepers against looters, anti-capitalist protestors against corporate interests no less than corporate interests against anti-capitalist protestors, burglars against householders no less than householders against burglars, paedophiles against neighbours no less than neighbours against paedophiles, and illegal immigrants against security guards sent to deport them no less than security guards against illegal immigrants.”³¹

These official duties are grounded in the interests of those whom the official seeks to protect: everyone within the relevant jurisdiction. (Thus the ‘community’ that officials protect is defined by nothing more than the shared presence in the territory over which the officials’ power extends. Anyone who enters that territory and is in need of protection has a claim on the officer’s attention and energy.) More generally, the rules governing the actions of officials are justified by the interests of all members of the community. This doesn’t mean that these rules must always effectively advance all of the interests of the community equally. But it does mean that they must be held to a standard demanding that they do, and that the right to make mistakes, and adopt rules that fall short of the standard, must itself be justified on the basis of the community members’ equally relevant interests.

This duty in turn explains why, when officials act on their reasonable but mistaken belief that someone poses a wrongful threat to others, and consequently

³¹ (Gardner 2013), p.111.

intentionally restrain or harm that innocent person, it is not unfair to ask the innocent victim to bear (part of) the costs of the mistake. In other words, it explains why officials like P are situated differently from citizens like C, whose situation we discussed in the previous Section: Because official P is acting on a general duty to protect everyone, a duty that is grounded in everyone's interest (including mine), I cannot complain that the official's special right not to suffer defensive harm for his mistakes gives him an unfair advantage, the way that S could normally have complained had C been given such a permission. Consider again the costs and benefits of such a permission that we discussed earlier: While the permission does enable P to avoid certain risks normally associated with reasonable mistakes, he does not thereby increase his discretion, since his exercise of the right to harm is not discretionary but subject to a duty. In particular, it does not allow P to err on the side of protecting himself and those he cares about at the expense of others, since his duty to intervene qua officer applies equally to whoever is wrongfully threatened, and whoever is doing the wrongful threatening. And while the risk of being arrested or restrained even though one is innocent does impose limits on one's choices, this burden is fairly distributed among all the members of the community, since the officer is again required to treat as equal every perceived offender.

This does not yet suffice to explain the special rights of officials. It merely shows that their special rights are themselves tied to their special duties; and once this link comes into focus, it might be asked why officials, but not ordinary citizens, have these duties as well as the corresponding rights. (It does explain why ordinary citizens who simply choose to protect others, even all others, may not have the corresponding rights: their choice is compatible with the absence of a duty, and thus leaves them with discretion that bears on how harm may be fairly distributed.)

There are two reasons for denying that a scheme assigning such duties to everyone would be morally justified. First, there is value in our having discretion to protect others (or not), and abandoning such discretion comes at a significant cost. The sheer fact that, if we abandoned the discretion, we may not have to bear the burden of our reasonable mistakes does not suffice to make a scheme that denies us such discretion attractive. In the case of officials, though the loss of discretion that they incur imposes a burden on them, that burden is justified, not because it saves them from the burdens of their reasonable mistakes, but because their duty in fact benefits the members of the community.

Second, the duties of officials that benefit the community are not exhausted by their obligation to protect everyone from wrongful threats. Officials do not simply have duties to do whatever they think would best protect innocent others. They also have duties of obedience, and are part of a hierarchy of decision-making. If each of us were simply required to act on what reasonably seems like a wrongful threat, then this would in fact significantly increase the risk of harmful conflicts.³² It is only because officials, though they have a duty to protect all citizens, also have a duty to accept as binding institutional decisions about how best to do so, that the duty to protect others in light of one's own reasonable judgments is not counterproductive. If P and Q are both officials, and they reach competing conclusions about whether S is liable to be arrested or punished, then this disagreement will itself be settled by the norms of the standing structure of authority that officials are part of. And as long as the officials' position in the hierarchy does not

³² Notice that here, unlike in the discussion of conflict in the previous section, the concern is not with competing permissions to take defensive steps, but with competing obligations. This significantly increases the risk of instrumentally harmful clashes.

depend on their arbitrary power advantages, neither does the resolution of disagreements among them. Consequently a scheme that imposes on officials both permissions and duties to act within the hierarchy does not run up against the worry that it imbues with moral relevance the morally arbitrary advantage of power that we earlier discussed as an important consideration against certain alternative schemes for distributing rights and duties among the members of our community.

Two features of this argument are worth highlighting. First, though the argument rests on appeals to the interests of all citizens, it is not a consequentialist argument in any straightforward sense. Most importantly, it is not an argument that simply asks what rules would maximize the overall or average wellbeing or safety of the community's members. The rules are justified by the benefits they provide. But the benefits they provide are sufficient to justify it only where they are distributed fairly; and fairness here requires, among other things, a concern with the avoidability of harm.

The first point is related to a second. Officials make arrests, I suggested earlier, based on judgments of liability. Though they may arrest or punish those who are innocent, they must do so only on the basis of a good faith judgment that they have, or are likely to have, committed a wrong. They may not (in general) arrest or punish someone because they judge that it would be socially optimal to do so. ("I know he is innocent, but I'll arrest him anyway to teach his friends a lesson.") One explanation for this is that it in fact reduces the risk we face to be harmed or restrained even though we are innocent. Harming the innocent is generally especially bad, because it is especially burdensome to suffer harm that we could not reasonably have avoided, and so inflicting such harm on us raises special concerns of fairness. But I think this is not all. We also care about how others treat us, and whether others recognize and treat us as someone who has

claims against being treated in certain ways. It is thus worse for us to be constrained or harmed by someone who does not even care whether we are liable to suffer such harm, than by someone who aims to inflict such harm on us only insofar as we are liable, but falls short of this ideal some time. This is especially true where the agent – individual or collective – acts in the name of the community as a whole, and thus expresses the community’s view of our status.

Let me conclude this Section by highlighting how the discussion here answers the libertarian challenge, to explain why public officials may do things – like arresting or punishing the innocent – that private persons must not do. The crucial point is that public officials have certain obligations, based on everyone’s interest, that private protectors lack. It is these obligations, and their grounding in the interests of everybody in the community, that explain why innocent persons mistakenly arrested or punished, though they are mistreated, lack the defensive and (some of their) compensatory rights that would normally arise from such mistreatment. The libertarian argument starts from the assumption that protection beyond that which we can provide as individuals must be acquired from private protection agencies who defend their clients against possible threats, but give no similar significance to the interests of non-clients in need of protection. (Perhaps the agency may find it practically simpler to protect clients and non-clients alike within a certain range of cases. But if so, the agency still does not provide protection for the non-client for their own sake, but only for the sake of minimizing costs to their clients. And so they cannot insist that the non-clients bear the cost of the agency’s mistakes, given that they treat protection of the non-clients as purely discretionary.) The agency is like the bodyguard: though it may have obligations to protect its clients, its non-clients cannot reasonably be expected to bear the burden of its mistakes.

5. Revisiting Officially Inflicted Harm

The account of officials' special moral situation offered in the previous sections explains the distinctive features of our moral practice regarding officially inflicted harm that, we saw earlier, standard views of the ethics of harm have a hard time accounting for. Let me start with the case of compensation, before turning again to the issue of defensive harm.

a. Compensation

If I am legitimately arrested, I will not normally be entitled to compensation even if I prove my innocence. If I am legitimately convicted and imprisoned, by contrast, I may be entitled to compensation when it turns out that I am in fact innocent. Yet in each case our rights against being non-consensually restrained seem to have been infringed. And indeed if an ordinary citizen like C engaged in either conduct, we may have compensatory claims against him. So these features of our practice regarding officially inflicted harm are puzzling.

They become, however, significantly less puzzling once we think of P's arresting S pursuant to his duty as part of a scheme that provides benefits to all the members of the community, and recognize that it is natural for the intended beneficiaries of a scheme to bear the costs of the mistakes that it makes. Consider a quite quotidian case that brings this out: A fiduciary administers the investment account of a child. The fiduciary satisfies all applicable duties of care and loyalty. In this case, the gains that result from his investment decisions will ultimately accrue to the child. So if the fiduciary's decision ends up not paying off, leading to a loss, it is only fair that those losses also fall on the account of the child for whose benefit the transaction was undertaken. A structurally similar point applies to officers making arrests: If the norms of arrest on which the officer acts are

grounded in the interests of all of the members of our community – justified by the benefit it provides to them – then everyone in the community ought to bear some of the burden of the officer’s mistakes.³³

This does not yet entail that compensation is necessarily inappropriate where the victim is also one of the intended beneficiaries. But it does entail that under certain circumstances, such compensation may indeed be unnecessary. This would most obviously hold true if the burdens were distributed fairly anyway, not just *ex ante*, but also *ex post*. (Imagine each of us would, over a life time, suffer roughly equal amounts of officially inflicted harm for innocent actions, and inequalities would be due only to choices we made for which we can reasonably be expected to bear the cost.) Yet this is quite clearly not the situation we find ourselves in: even if the risk of being mistakenly arrested were equally distributed *ex ante*, they would not fall equally on all of us *ex post*. But even in this situation, of *ex post* inequality, there may be good reasons not to compensate those of bear greater costs than others. For if the cost of administering such compensation are quite high, and must be borne by the beneficiaries of the scheme, then the would-be recipients of such compensation may still find it counterproductive to establish a scheme providing such compensation.³⁴

³³ For a related thought about collective responsibility, see (Stilz 2011). Stilz, however, adds excessively demanding conditions for liability, according to which the citizens’ will must be implicated in the actions of their community (either via democratic decision-making or identification). By contrast, on the view sketched here, the focus is simply on whether the citizens are both *de facto* beneficiaries of the scheme and those whose interests justify that scheme in the first place.

³⁴ This is not to deny that there may also be circumstances where compensation for arresting an innocent person based on reasonable suspicion is appropriate. In societies like ours, people regularly become suspect based on the color of their skin or ethnic heritage. Often this reflects outright racism. But at times it also reflects the fact that people tend to be less reliable in distinguishing the appearances of members of an ethnic group other than their own. (For a detailed survey and summary of empirical studies to this effect, see (Meissner and Brigham 2001).) As a result, the suspect is identified more generically, and innocent members of the group are more likely to fall under the description offered. Compensation may be appropriate then, not just for the cost the arrest imposes as such, but also for the particular indignity of being arrested based on another’s reliably unreliable judgment of one’s appearance.

Yet this is much less plausible where an innocent suspect has been convicted to a lengthy prison sentence. Even if the ex ante distribution of this risk is fair, it is fairer still to ensure ex post that those who are not liable for imprisonment do not bear the full burden of the mistakes the scheme inevitably makes, even if such compensation comes at some administrative cost to all of us. More generally, we can explain why some harms inflicted by someone acting for the innocent victim's benefit give rise to compensation, and others do not, by identifying the relevant benefits and the fair distribution of their correlative burdens as part of a larger scheme providing services and goods to the community as a whole.

b. Defensive Harm

Let me next return to the issue of defensive harm. P, we have seen, is not expected to have to suffer defensive harm at the hands of an innocent person he arrests or punishes as long as he acts on the basis of a reasonable belief that the person is liable to be arrested or punished. The discussion, in the previous section, of the neighbor who protects the abandoned house and mistakenly restrains its owner provides the blueprint for explaining why officials can avoid the risks associated with defensive harm that would normally flow from their reasonable mistakes: because the duty on which the officer acts is justified, at least in part, by the interests of the person who ends up being harmed by the officer's action, and the officer is already taking risks to protect these interests, it would be unfair to also ask the officer to bear the additional burden of suffering defensive harm in cases of reasonable mistakes.

Notice that, though in the case of the neighbor discussed earlier, I assumed that it was S's ownership interest in the house that grounded the neighbor's duty to protect the house against trespassers, it need not be S's interest alone. Imagine that the house is

owned, not solely by S, but by S and her nine siblings. The neighbor's duty of protection rests on the ten siblings' interest in their property; and yet I think intuitively S would still have to bear, in full, the burden of the neighbor's reasonable mistake where he restrains her for apparent trespass. (She may have to be compensated later on by her siblings. But where the only question is how to distribute harm between her and the neighbor, she must accept that it falls on her.)

Consider another case, that may bring this out even more strongly:³⁵ S and nine other people have been kidnapped by K, who threatens to kill them. P is an officer who is sent to rescue them. To confuse the police, K has put all the victims and herself in identical clothes. But, as the victims have been able to communicate to the police, only K has a gun, and can be identified on this basis. Just as P begins his mission, S, by a stroke of luck, manages to overpower K and take the gun. As P enters the building, he sees S, and reasonably takes her to be K. He lifts his gun to shoot her. S recognizes that P believes she is K. Her only way of stopping P from shooting her is to shoot P first. I am strongly inclined to think that here, S must not shoot P, even if this comes at the cost of being shot by him.

These cases also help sharpen the conditions that must be met if S can reasonably be asked to bear the burden of P's mistake. To begin, the issue is not P's motivation: When the commander sought for someone to fulfill the mission to free the kidnapped, P may have volunteered because his friend was among the victims. He wouldn't have volunteered to save S. Still, once he has volunteered, he has a duty to save all the victims; and where he can only save some, he cannot give preference to his friend over other

³⁵ This is a variation on Tadros' 'Soldier Rescue' case: (Tadros 2011), p.233.

victims. If, by contrast, P's only duty was to rescue his friend, and saving the others were merely a coincidental benefit, I doubt that S would have to bear the same burden of P's mistake.

Notice too that, though the relevant duty must be to S's benefit, it need not be the case that S's interest by itself is sufficient to justify that duty. To see this, return to the case of the helpful neighbor. Imagine the neighbor has the duty to stop trespassers only because it both benefits the home owner (S and her siblings) and others who live on the street, and wish to prevent squatters from moving into their neighborhood. Neither interest by itself would be sufficient to ground the neighbor's duty. But the interests are jointly sufficient, and individually necessary (in this case), to ground a duty of neighborliness to look after the house. In that case, S and the other neighbors who benefit from avoiding squatting would have to bear the burden of mistakes made by the good neighbor's fulfilling the duty.

Similarly, P's duty to protect everyone within the jurisdiction from wrongful threats need not rest solely on S's interest for it to be the case that S can be asked to bear the burden of P's occasional mistake. This demand can be made on S as long as her interest is part of the considerations that justify P's duty. If it does play this role, then this will, for instance, put pressure on P (or P's superiors) to adjust the duty in light of the effects that it has on the interests of S. If the existing system of rules is, for instance, predictably imposing an unfair burden on S, then this is a reason for the system to be revised; and if, in the face of such unfairness, the scheme is not revised, it is revealed as not being genuinely grounded in the interests of S (and others like her).

This last point has important implications that I can only flag here, but that are worth developing in more detail elsewhere: If a system predictably imposes unfair

burdens on some of those to whom it regularly applies, and if no effort is made to correct for this unfairness, then those who are treated unfairly in this way cannot be asked to bear the full burdens of the cost of the system. It is not just that they lack reason to trust the system, and thus to do their part where they are asked to. They may also reasonably deny that they should have to bear the burden of mistakes made by officials following rules that do not in fact seem to be responsive to their interests. If the criminal justice system appears to function to protect them, and control us, then we need not bear the burdens of the mistakes it makes, and police officers really are more like soldiers in occupied territory than servants of the community they patrol – with all of the consequences this has for their liability to suffer defensive harm by innocent victims of their mistakes.

Bibliography

- Alexander, M. (2010). The New Jim Crow: Mass Incarceration in the Age of Colorblindness. New York, The New Press.
- Applbaum, A. I. (2000). Ethics for Adversaries: The Morality of Roles in Public and Professional Life. Princeton, Princeton University Press.
- Bennett, C. (2011). "Expressive Punishment and Political Authority." Ohio State Journal of Criminal Law **8**: 285-318.
- Duff, R. A. (2001). Punishment, Communication, and Community. Oxford, Oxford University Press.
- Dworkin, R. (1985). Principle, Policy, Procedure. A Matter of Principle. Cambridge (MA), Harvard University Press: 72-103.
- Estlund, D. (2007). "On Following Orders in an Unjust War." The Journal of Political Philosophy **15**(2): 213-234.
- Estlund, D. (2008). Democratic authority: a philosophical framework. Princeton, Princeton University Press.
- Feinberg, J. (1978). "Voluntary Euthanasia and the Inalienable Right to Life." Philosophy and Public Affairs **7**(2): 93-123.
- Firth, J. M. and J. Quong (2012). "Necessity, Moral Liability, and Defensive Harm." Law and Philosophy **31**(6): 673-701.
- Gardner, J. (2013). Criminals in Uniform. The Constitution of Criminal Law. R. A. Duff, L. Farmer, S. E. Marshall, M. Renzo and V. Tadros. Oxford, Oxford University Press: 97-118.
- Greenawalt, K. (1998). "Justifications, Excuses, and a Model Penal Code for Democratic Societies." Criminal Justice Ethics **17**: 14-28.
- Harel, A. (2008). "Why Only the State May Inflict Criminal Sanctions: The Case Against Privately Inflicted Sanctions." Legal Theory **14**: 113-133.
- Huemer, M. (2013). The problem of political authority : an examination of the right to coerce and the duty to obey. Houndmills, Basingstoke, Hampshire ; New York, Palgrave Macmillan.
- McMahan, J. (2009). Killing in War. Oxford, Oxford University Press.
- Meissner, C. A. and J. C. Brigham (2001). "Thirty Years of Investigating the Own-Race Bias in Memory for Faces: A Meta-Analytic Review." Psychology, Public Policy, and Law **7**(1): 3-35.

Quong, J. (2012). "Liability to Defensive Harm." Philosophy and Public Affairs **40**(1): 45-77.

Renzo, M. (2013). "Democratic authority and the duty to fight unjust wars." Analysis **73**(4): 668-676.

Ripstein, A. (2009). Force and Freedom: Kant's Legal and Political Philosophy. Cambridge (MA), Harvard University Press.

Rodin, D. (2011). "Justifying Harm." Ethics **122**(1): 74-110.

Simmons, A. J. (1979). Moral principles and political obligations. Princeton, N.J., Princeton University Press.

Simmons, A. J. (2001). Justification and legitimacy : essays on rights and obligations. Cambridge ; New York, Cambridge University Press.

Stilz, A. (2009). Liberal Loyalty: Freedom, Obligation, and the State. Princeton, Princeton University Press.

Stilz, A. (2011). "Collective Responsibility and the State." The Journal of Political Philosophy **19**(2): 190-208.

Tadros, V. (2011). The Ends of Harm. Oxford, Oxford University Press.

Tadros, V. (2012). "Duty and Liability." Utilitas **24**(2): 259-277.

Thomson, J. J. (1990). The realm of rights. Cambridge, Mass., Harvard University Press.

Thorburn, M. (2008). "Justifications, Powers, and Authority." Yale Law Journal **117**: 1070-1130.

Viehoff, D. (2014). "Democratic Equality and Political Authority." Philosophy and Public Affairs **42**(4).

Viehoff, D. (2016). "Authority and Expertise." The Journal of Political Philosophy **24**(4): 406-426.