The Price of Duty

Abstract

State officials regularly impose harm on the citizens they are supposed to serve, some of it wrongful. Who should bear the burden of paying compensation for these wrongs? Should it be the agents themselves, or should the burden be spread across the citizenry via taxation? This essay develops a theory of *limited official immunity*, according to which citizens have a moral duty to assume the costs of (only) certain official wrongs. I argue that limited official immunity is an upshot of a general principle of distributive justice, according to which those who are morally required to participate in some project are required to share certain costs of that project. Understanding official immunity as derivative of this principle helps us to identify, not just the grounds, but also the limits of official immunity: it helps us to sort between the kinds of official wrongs for which the burdens should fall on the individual agents themselves and the kinds of wrongs for which those burdens should be assumed by the citizenry.

Keywords: beneficiary pays, compensation, distributive justice, immunity, responsibility

1 Introduction

State officials regularly impose harm in the line of duty, some of it wrongful.[1] For many of these wrongful harms, however, it is widely believed that state officials should not have to pay compensation. Consider, for example:

> *Fire*. Citizen's house is hit by lightning and erupts in flames. Citizen is trapped inside. To save Citizen, Firefighter blasts the house with water, extinguishing the flames. As an unintended and unavoidable side-effect, she causes thousands of dollars of water damage to Neighbor's home.

> *Mistaken Imprisonment*. Judge exercises great care in his handling of a criminal case, and accurately applies the law in light of his available evidence. He sentences Victim to prison. Guard transports Victim to prison and spends many years preventing Victim from escaping — until new evidence emerges proving that Victim was innocent all along.

Neighbor and Victim suffer wrongful harms for which they are morally owed compensation. But it seems clear that Firefighter should not have to be the one to compensate Neighbor, and that neither Judge nor Guard should have to be the one to compensate Victim. Compensation should be paid — perhaps by the state, perhaps by the citizenry — but not by these individual officials.

Let's give the general phenomenon a name. Not to be confused with the legal phenomenon of the same name, we'll call it

[1] By 'wrongful harm' I mean harm that is rights-infringing. Because rights may sometimes be permissibly infringed, a wrongful harm (on this way of using the term) needn't be morally impermissible. Likewise, the agent of wrongful harm needn't be blameworthy.

> **Official Immunity**. State officials (morally) should not have to bear the costs for some of the wrongful harms they impose in the line of duty.

This way of characterizing the phenomenon is meant to be ecumenical between various forms this moral immunity might take. It might be, for example, that Firefighter's immunity consists in the simple absence of a moral duty to compensate Neighbor. Or it might be that Firefighter does have a pro tanto duty to compensate Neighbor, but that she *also* has a right of *indemnification* — that is, a right that some third-party (such as the state) assume the costs of her compensatory duty. As I'll use terms, either of these would constitute a form of official immunity.

The present essay develops a theory of the grounds, nature, and limits of official immunity. Why do state officials enjoy such immunity? What form does this immunity take? What are the sorts of harms for which state officials do (not) enjoy immunity?

The essay proceeds as follows. Sections §2-§4 canvas a number of potential explanations for official immunity that I find wanting. §2 considers various ways of explaining official immunity by appeal to the *principal-agent relationship* that holds between the state and its officials. §3 marks a turn from looking at the relationship between the state and its officials to looking at the relationship between the state and the citizenry. §3 and §4 consider and reject different versions of the idea that official immunity is explained by the fact that citizens benefit from the actions of state officials. §5 introduces my positive account — the Price of Duty account — according to which official immunity is explained by the special duties of cost sharing that fall on persons for whom participation in a cooperative project is morally required. §6 shows how this account helps us make progress in understanding the *limits* of official immunity. Specifically, it helps us identify four types of damages for which state officials are *not* immune. §7 concludes by noting the justificatory burden this account

puts on the regime of legal immunity that we find, for example, in the United States.

2 The Principal-Agent Relation

Here is a notable feature of many cases of official immunity: a private citizen acting on their own *would* be morally required to pay compensation for many of the sorts of wrongful harms for which state officials are not required to pay compensation. Private citizens are often morally liable where state officials are morally immunie. Consider, for example, a variant of the Fire case:

> *Fire 2*. Citizen's house is hit by lightning and erupts in flames. There isn't time to wait for the fire department, and Citizen can save himself only by putting the fire out with his garden hose. As an unintended and unavoidable side-effect, he causes thousands of dollars of water damage to Neighbor's home.

Whereas it seems Firefighter should not have to compensate Neighbor in the original Fire case, it seems that Citizen *should* have to compensate Neighbor in Fire 2. What's the relevant difference? One candidate that jumps out is the fact that Firefighter acts as an *agent* of another (the state or the citizenry) whereas Citizen acts under his own authority. This difference is especially notable in light of the fact that the principal-agent relationship has been long recognized as a source of immunity in private law in many nations. If, for example, an employee of trucking company A damages some of the cargo she is hauling for company B, the employee may enjoy a legal right to indemnification against her employer. That is, although the employee would have a legal duty to compensate company B for the damages she caused, company A would have a legal duty to cover those compensatory costs.[2] (This isn't the sort of immunity that

[2] American Law Institute, *Restatement (Third) of Agency* (St. Paul, Minnesota: American Law Institute Publishers, 2006), ch. 7.

consists in the simple absence of duty; rather it's the sort of immunity that consists in a right that someone else bear the costs of one's duty — a right of indemnification.)

Now it might be that there is no justification for this feature of private law. Or it might be that it admits only of an instrumentalist justification: perhaps this policy does not serve to enforce any antecedent moral duties that employers have to indemnify their employees, but is justified only by, for example, its distributive or incentive effects.[3]

Perhaps. In this section, however, I want to take seriously the idea that there is a deeper moral justification behind private agency law. I want to take seriously the idea that agents have an antecedent moral right to be indemnified by their principals for certain costs they cause in their capacity as agents. I consider two stories that might be told about why agents have such rights. I argue, however, that these explanations cannot plausibly extend to account for the moral immunity enjoyed by state officials.

2.1 The Promissory Account

One story claims that principals have a moral duty to indemnify their agents because they have *promised* to do so. In the case of private employers, for example, it might be thought that they make a promise to indemnify their employees — perhaps not explicitly, but at least implicitly, in virtue of the fact that it is conventionally understood that employers assume a duty of indemnification when they hire someone. Perhaps the same is true of state officials: they have a right to be indemnified by either the state or the citizenry when and because the state or the citizenry has

[3] A distributive argument for this policy might appeal to the fact that employers are typically better positioned than their employees to spread costs (to consumers, insurers, and the like). An incentive-based argument might appeal to the fact that employers — having the ability to make organizational changes — are typically better positioned than their employees to reduce future employment-related damages. Since an indemnification requirement would incentivize employers to make such changes, total future damages would be minimized.

promised to indemnify them. Call this the **Promissory Account**. On this account, promises are the grounds of official immunity.

Whether or not promises explain the immunity enjoyed by private employees, we should reject Promissory. Consider, first, the idea that state officials are immune when and because the *state* has promised to indemnify them. The problem with this idea is that official immunity is not plausibly contingent on whether the state has promised to indemnify its officials. Imagine a state that has *not* promised to indemnify its officials. That is, imagine a state whose law does not bestow on the state a legal duty to indemnify its officials. But now consider our paradigm cases of official immunity against the backdrop of such a state: Firefighter causes damage to Neighbor's home; Judge and Guard mistakenly imprison the innocent Victim. Should these costs fall on Firefighter, Judge, or Guard in a world where the law is silent about indemnification? Surely not. Altering the legal backdrop in this way does not change the intuitive judgment that Firefighter, Judge, and Guard should not — morally speaking —- have to bear the costs of compensating Neighbor or Victim.

The same problem holds for the idea that state officials are immune when and because the *citizenry* has promised to indemnify them. And here we needn't even use our imaginations: it isn't plausible that most citizens in *real-world* political communities promise to indemnify their officials. Promises are complicated creatures, but the following is a plausible minimal constraint on the making of a promise:

> **Epistemic Constraint**. X promises to ϕ only if X performs some action A such that (i) X believes or intends that A constitutes a promise to ψ (where ψ-ing is or entails ϕ-ing), or (ii) it is public knowledge that A is a type of action that constitutes a promise to ψ.

This principle is extremely plausible. Promises alter the moral boundaries

of the promisor: they alter their duties and liabilities. Since it is in our interest to have control over such changes to our moral boundaries, it is in our interest to have a good amount of control over when we make promises. Likewise, because promises also change the moral boundaries of the promisee (who takes on new rights) and because social coordination often requires us to understand our moral relations to one another, it is also in our interest to be able to know when *others* make promises. But we'd have neither much control over when we make promises, nor find it easy to know when others make promises, if Epistemic Constraint were false.[4]

This plausible principle makes serious trouble for the idea that most citizens have promised to indemnify their state officials. If asked, most citizens would surely assert that they have never promised such a thing. This is good evidence, not just that those citizens haven't done something they believed or intended to constitute a promise to indemnify, but that they haven't even done something that is publicly known to constitute a promise to indemnify. It is very implausible that it could be public knowledge among the citizenry that ϕ-ing constitutes a promise to indemnify state officials but that most citizens have done ϕ without believing that they have promised to indemnify state officials.

2.2 The Responsibility Account

The Promissory account isn't promising. But here's a different way one might think to explain official immunity in terms of the principal-agent relation. Perhaps the principal-agent relationship is the source of official immunity, not because it involves certain promises, but because of its effect on the distribution of responsibility.

The idea is this. Compensatory duty tends to track responsibility. The

[4]For an excellent discussion along these lines, see Renée Jorgensen, "Moral Risk and Communicating Consent," *Philosophy and Public Affairs* 47, no. 2 (2019):179-207. Jorgensen is focused on consent, rather than promise, but much of her discussion applies to both phenomena.

party who typically owes the most compensation for some harm is the party who was most responsible for producing the harm.[5] But agents are *instruments* of their principals. As such, when agents cause wrongful harm in their capacity as agents, it is the principal who is the "first cause", who is ultimately in control, and who is thus the primary locus of responsibility for (many of) the wrongful harms produced by the agent. Because compensatory duty tracks responsibility, then, it is the principal, rather than the agent, who ought to shoulder the burden of compensation for these wrongs.

Call this the **Responsibility Account** of official immunity: state officials should not have to pay compensation for a wrongful harm they produce when and because the state (or some state institution) bears the bulk of responsibility for the production of that harm.[6] On this account, the salient difference between Firefighter damaging Neighbor's home (Fire) and Citizen damaging Neighbor's home (Fire 2) is that Firefighter is not chiefly responsible for the damage she does (it's the state that is most responsible) whereas Citizen *is* chiefly responsible for the damage he

[5]The word 'responsibility' is polysemous, having many different, closely-related meanings. There is one sense of 'responsibility' on which it is tautologous that the person who bears most of the responsibility for a harmful wrong is the person who has a duty to pay compensation for that wrong. Some philosophers have called this *duty responsibility*. (See, for example, Gary Watson, "Raz on Responsibility," *Criminal Law and Philosophy* 10, no. 3 (2016): 295-409.) *Part of what it is* to be duty responsible for harm H is to have a duty to provide compensation for H. The sort of responsibility I'm referring to in this paper is different. Call it *agent responsibility*. Agent responsibility for H consists in a form of control between one's agency and H; the tighter this connection, the greater one's agent responsibility. Compensatory duty for H is not part of what it is to be agent responsible for H, though the fact that someone has a compensatory duty for H is often *explained by* the fact that they are agent responsible for that harm. A similar distinction is made by Joseph Raz, "Responsibility and the Negligence Standard," *Oxford Journal of Legal Studies*, 30 no. 1 (2010): 1-18.

[6]A more extreme account in the neighborhood is what we might call the Vicarious Agency Account. On this account, it isn't just that the state bears the bulk of responsibility for the harm in question, but that it is the state (rather than the official) that is the *agent* of the harm; the official is merely a conduit of state agency. I won't spend any time discussing this account, however, as it has less prima facie plausibility than the Responsibility Account and holds up even worse against the objections I press in this section.

does.

Some might object to the idea that the balance of responsibility can come apart from the balance of causal contribution. But this clearly happens on occasion. Consider, for example:

> *Overboard*. Reckless is driving his new speedboat on a small lake. Wanting to see how fast his boat can go, he presses the throttle all the way forward, bringing his speed to more than twice the legal speed limit. At such high speeds he fails to even notice when his boat crashes into Paddler's canoe, destroying the canoe and throwing Paddler into the water. Paddler can save herself from drowning only by climbing onto a nearby dock. But she can only do so by pulling herself up by the finger of Sleepy, who is sunbathing at the end of the dock (and who cannot be woken to give his consent). Paddler pulls herself up, breaking Sleepy's finger.

Sleepy is owed compensation for his broken finger. But although it is Paddler who chose to injure Sleepy and who was the direct cause of Sleepy's injuries, the burden of compensation surely falls on Reckless. This, we want to say, is because Reckless is the person who bears chief responsibility for Sleepy's injury.

The problem with Responsibility isn't that it makes a conceptual distinction between causal contribution and responsibility. As the above example shows, that's an important distinction. The problem is that the account undersells the extent to which state officials are responsible for the harms they produce. For one thing, state officials are *not* mere instruments of the state; they do not merely "run the code" written by their superiors. They often enjoy considerable discretion as to *how* they follow state directives. Some harms imposed by state officials are *non-discretionary* (in legal parlance, "ministerial"): these are harms that cannot

feasibly be avoided so long as the official follows the state's directives. But many harms are *discretionary*: these are harms that are contingent on the official choosing one way of following state directives rather than another. For many discretionary harms, state officials exercise much more control over those outcomes than their principals, and as such are more responsible for producing those harms than their principals. And yet the line between official immunity and official non-immunity is surely not the line between discretionary and non-discretionary harms: some discretionary harms are among the paradigm harms for which officials enjoy compensatory immunity. Recall the case of Mistaken Imprisonment. Here, we noted, it seems clear that Judge should not have to compensate Victim for his wrongful imprisonment. But suppose we were to learn that Judge was not legally required to give any particular sentence, and that it was left to his discretion to sentence Victim to anything between five and ten years in prison. As it happened, Judge sentenced Victim to seven years in prison. Does this negate Judge's immunity? Surely not. The intuitive verdict that Judge should not have to compensate Victim is not altered by learning that Judge exercised discretion in the sentencing.

Here's another reason we should reject Responsibility. Reflect on our attitudes toward the *benefits* produced by state officials. Imagine that Firefighter is ordered on another mission. She is told to parachute into a remote town to slow the spread of an approaching forest fire. She performs the mission exactly as directed and her efforts save the lives of dozens of people. Although Firefighter saved the townsfolk because she was simply following orders, most of us would think it fitting to heap praise and honors on Firefighter. We would praise her, not just for being good at her job, but especially for the good outcome she brought about by being good at her job; we would praise her for saving the lives of the townsfolk. Moreover, we would think that Firefighter deserves *greater* praise for saving those people than does the state itself. This is no doubt because we judge Jumper to be the agent who is *most* responsible for the

rescue. But if state officials bear chief responsibility for the benefits they produce by their obedience, it is very hard to see why they would not bear chief responsibility for the harms they impose by their obedience as well.[7]

3 The Beneficial State Account

The Promissory and Responsibility accounts suffer from serious problems. As I can see no other way in which the principal-agent relation might explain official immunity, I now turn to explore elsewhere. In the next three sections I explore the idea that official immunity is grounded, not in a relation between officials and the state, but instead in a relation between officials and the *citizenry*.

Many philosophers have defended the so-called *Beneficiary Pays Principle*. There are different versions of this principle, but the common core is that

> **Beneficiary Pays**. The receipt of benefits resulting from wrongful harm can sometimes ground a duty on the part of the beneficiary to provide compensation to the victim of that harm.[8]

[7]See Jeff McMahan, Killing in War (Oxford: Oxford University Press, 2009), 84-91.

[8]Different variants of this idea are defended by Christian Barry and Robert Kirby, "Scepticism about Beneficiary Pays: A Critique," Journal of Applied Philosophy 34 no. 3 (2017): 282-300; Saba Bazargan-Forward, "Grounding the Beneficiary Pays Principle," Oxford Studies in Political Philosophy 8 (forthcoming); Daniel Butt, "'A Doctrine Quite New and Altogether Untenable': Defending the Beneficiary Pays Principle," Journal of Applied Philosophy 31 no. 4 (2014): 336-348; Alexandra Cuoto, "The Beneficiary Pays Principle and Strict Liability," Philosophical Studies 175 no. 9 (2018): 2169-2189; Robert Goodin, "Disgorging the Fruits of Historical Wrongdoing," American Political Science Review 107 no.3 (2013): 478-491; Sigurd Lindstad, "Beneficiary Pays and Respect for Autonomy," Social Theory and Practice 47 no. 1 (2012): 153-169; Tom Parr, "The Moral Taintedness of Benefiting from Injustice," Ethical Theory and Moral Practice 19 no. 4 (2016): 985-996; Victor Tadros, "Orwell's Battle with Brittain: Vicarious Liability for Unjust Aggression," *Philosophy and Public Affairs* 42, no. 1 (2014):42-77; Judith Thomson, "Preferential Hiring," Philosophy and Public Affairs 2 no. 4 (1973): 364-384; Daniel Viehoff, "Legitimacy as a Right to Err," in NOMOS LXI: Political Legitimacy, eds. Jack Knight and Melissa Schwartzberg (New York: NYU Press, 2019), 174-200; and Viehoff, "Legitimate Injustice and Acting for Others," Philosophy and Public Affairs 50 (2022): 301-374.

Defenders of Beneficiary Pays typically claim that the beneficiary's agency needn't bear any connection to the wrong or to their own benefit. The beneficiary need not have caused or contributed to the wrong. The beneficiary need not have asked for or intended the benefit. No authorization is needed. The beneficiary need only *receive* the benefit.

It's natural to think to explain official immunity in terms of something like Beneficiary Pays. As formulated, however, the principle plainly won't do the job. The principle focuses on the beneficiaries of *particular* instances of wrongful harm. But there are many instances where state officials cause wrongful harm without benefiting anyone — instances where the official seems nonetheless immune from owing compensation. No one benefits from the actions of Judge or Guard in Mistaken Imprisonment, for instance. And we can easily imagine a variant of Fire in which no one benefits from Firefighter's action: perhaps she isn't able to extinguish the fire despite all the water she blasts on the house. The absence of benefit in these instances doesn't matter: whether Judge, Guard, or Firefighter should have to pay compensation for the harms they impose is not contingent on whether those particular harms benefit anyone.

The upshot is that if we are going to appeal to the benefits of the actions of state officials to explain official immunity, we need to appeal, not to the benefits that result from *particular* harms, but rather to the benefits that result from some *broader project* that gives rise to the harms in question. The idea is this. When someone produces a wrongful harm in the execution of a project from which the Xs benefit, the Xs thereby incur duties to bear (at least some of) the burden of compensating for that harm. The state's project of maintaining a just and good political order is a project from which its citizens generally benefit. This — the story goes — is why the citizenry is duty-bound to shoulder the costs of certain official wrongs. Call this the **Beneficial State Account**.[9] On this account, the crucial

[9]It should be noted that many versions of the Beneficiary Pays Principle will not get us anything like this account. For example, some theorists (e.g., Bazargan-Forward,

difference between Fire and Fire 2 is the fact that Firefighter's actions (in Fire) are part of a broader project that serves the citizenry, whereas Citizen's actions (in Fire 2) are not.

Let's interrogate the idea at the heart of this account — the idea that duties of compensation can fall on the beneficiary of a wrongful harm simply *because* they are a beneficiary. This idea is clearly implausible, absent some kind of restriction. To see why, consider a case like:

> *Three Fishermen*. There are three fishermen (Al, Bob, and Chuck) that share the shoreline of a bay and compete for business. In the dead of night and in a drunken fit of jealousy, Al sneaks onto Chuck's dock and burns Chuck's boat. While making his escape, however, Al slips and breaks his arm. As a consequence of these events, neither Al nor Chuck are able to fish for many weeks. This helps Bob to catch more fish and make a greater profit than he otherwise would have.

Two things are clear. The first is that Al has a duty to compensate Chuck for destroying his boat. The second is that Bob — the only beneficiary of Chuck's loss — has no duty to assume Al's compensatory burden. Neither Al nor Chuck can legitimately demand this of Bob.

Compensatory duty clearly requires more than the mere receipt of benefit. If the Beneficiary Pays idea is to be at all plausible, we need some way of restricting the idea to a certain subset of beneficiaries. Can we locate a plausible restriction?

One possibile restriction takes its inspiration from a famous quote from Rawls regarding duties of "fair play":

Goodin, and Parr) defend only a *disgorgement* version of the principle, according to which beneficiaries of a harm needn't contribute to offsetting the victim's losses, but need only transfer their own *gains* to the victim.

> When a number of persons engage in a mutually advantageous cooperative venture according to rules, and thus restrict their liberty in ways necessary to yield advantages for all, those who have submitted to these restrictions have a right to similar acquiescence on the part of those who have benefited from their submission.[10]

Rawls is talking about duties that beneficiaries have to bear the costs of *obedience*, not the costs of compensation. But we might think that his way of restricting this principle is also a nice way of restricting the Beneficiary Pays idea. Perhaps there is something special about *cooperative* projects; perhaps one does not acquire duties of assumption by benefiting from just any wrong, but only from wrongs that are the products of cooperative projects.

The cooperative-project restriction doesn't help, however. There are all sorts of cooperative projects for which beneficiaries plainly do not have duties of compensatory assumption. Consider, for example:

> *Bells*. A small minority of people in your neighborhood have installed large church bells on their properties, and have been working together over the past few months to produce beautiful music at the top of every hour, to the enjoyment of everyone in the neighborhood. One night, however, while one of the bell-ringers is working the bells, the rope frays. The bell is about to fall on Bell Ringer. He avoids being crushed in the only way he can — by redirecting the bell onto his neighbor's unoccupied car.[11]

[10]John Rawls, *A Theory of Justice* (Cambridge, Massachusetts: Harvard University Press, 1971), 108-114. We get a very similar principle from H.L.A. Hart, "Are There Any Natural Rights?" *Philosophical Review* 64 no. 2 (1955): 185.

[11]This case is a bow to Robert Nozick's case of the neighborhood PA system from *Anarchy, State, and Utopia* (Wiley-Blackwell 2001): 93-95.

Bell Ringer clearly incurs a duty to compensate his neighbor for the damage to her car. Does everyone in the neighborhood who has benefited from the tolling of the bells owe it to Bell Ringer to offset his duty? Surely not. (In §5 I'll explore the relevant difference between state officials and the bell ringers.)

Here is a different restriction we might try. We might depart from the standard emphasis on benefits that are *involuntarily* received, and insist that beneficiaries take on duties of assumption only when they have in some sense *accepted* the benefits.[12]

But this move won't work either. In Three Fishermen, for example, Al's wrong makes it so that Bob has the opportunity to catch more fish. But Bob receives the benefits of this opportunity only by going out on the water and intentionally hauling in the extra fish. The benefit is not dropped in his lap; he goes out and claims it. What's more, this benefit is such that Bob could easily avoid it. We can imagine that it would be no trouble to avoid catching extra fish. And yet despite all this Bob does not have a duty to offset Al's compensatory burden.

If anything, the sense in which citizens "accept" the benefits of the state is *less* robust than the sense in which Bob accepts his benefits. It is easy for Bob to avoid catching extra fish. It is not easy to avoid the benefits that are conferred on citizens by the state. Most of these benefits can be avoided only at tremendous cost — for example, emigration, social isolation, or poverty.[13]

4 Agency on Loan

The most developed and creative defense of a variation of the Beneficiary Pays idea comes courtesy of Daniel Viehoff, who has explicitly developed

[12]See A. John Simmons, "The Principle of Fair Play," *Philosophy and Public Affairs* 8 no. 4 (1979):307-337 for an attempt to formulate the principle of fair play in terms of accepted benefits.

[13]C.f., Simmons, "Fair Play," 334-335.

a theory of official immunity on the grounds of this idea.[14] His distinctive account merits a section unto itself.

Viehoff's account starts with the basic and uncontroversial idea that a person typically must shoulder the costs that their own resources impose on others. If my billy goat manages to escape the fence I've constructed and devours my neighbor's garden, the costs of those damages are mine to bear.

We can, however, *lend* our resources to others. And when we do so, it the person to whom we lend our resources who must bear the costs those resources impose on others. If I lend you my billy goat for a week because your lawn mower is broken, and he escapes your custody, the costs of his mischief are yours to bear rather than mine. This is because, by lending you my goat, he "becomes, at least as between you and me, temporarily your resource rather than mine" (329).

Viehoff argues that it's not just our external resources that we can loan out to others. We can also loan out, in restricted ways, our body and our agency. Such lending, he claims, can often be seen in cases where one person acts as a surrogate decision maker for someone who is unconscious and unable to make their own medical decisions. The surrogate, provided they satisfy certain conditions, effectively lends their agential capacities to the patient, making those capacities "asymmetrically available for the pursuit of [the patient's] interests rather than [their own]" (331). When the surrogate succeeds in lending their agency to the patient in this way, the surrogate is immune from owing compensation for mistakes they make as a surrogate.

One person, A, loans out their agency to B on some matter, Viehoff claims,

[14]This theory is first developed in Daniel Viehoff, "Legitimacy as a Right to Err," in *NOMOS LXI: Political Legitimacy*, eds. Jack Knight and Melissa Schwartzberg (NYU Press, 2019), 174-200, and more fully developed in Viehoff, "Legitimate Injustice and Acting for Others," *Philosophy & Public Affairs* 50 (2022):301-374.

just in case two conditions are satisfied. The first condition concerns the way in which the A exercises her agency on B's behalf:

> **Deliberative Condition**: A loans her agency to B on some matter only if her choices in the matter are guided by B's interests (and by the interests of third-parties to the extent that it would be reasonable for B, himself, to be guided by those interests).[15]

The deliberative condition requires that A treat her own agency as a resource temporarily assigned to B.[16] But this isn't enough for A to succeed in putting her agency on loan to B. After all, B "may have reasonable grounds for objecting to [A's] agency becoming one of his resources" (336). For this reason, Viehoff posits a second necessary condition:

> **Justificatory Condition**. A loans her agency to B on some matter only if either (i) B has given his actual consent to A's granting her agency to B as a resource of his, or (ii) it is in B's "permissibly pursuable personal interest" *in expectation* that A grant her agency to B as a resource of his.[17]

[15]There are various complications that this gloss of the Deliberative Condition leaves out, but they are ones we can safely ignore for our purposes. See pages 332-336 of "Legitimate Injustice" for a detailed discussion of this condition. See also his discussion in "Legitimacy," 192.

[16]Ibid., 336

[17]Ibid., 337. The "in expectation" bit is an important qualification. Viehoff does not think that condition (ii) requires that B is made better off, *ex post*, by the lending of A's agency. He says:

> When Ara acts for Pete, Pete incurs liability even if Ara makes a mistake, and chooses to do what she ought not to have done. This may be so even if what Ara does while acting for Pete is objectively unreasonable, as long as both the [Deliberative and Justificatory conditions] are met. Why? Because Pete has no general entitlement that Ara possess the capacity to avoid these mistakes, nor that she make reasonable choices only. He isn't, that is, generally entitled that she be able to pursue *his* personal projects and goals with

When both the Deliberative and Justificatory conditions are met, A succeeds in loaning her agency out to B. When she does so, her agency becomes partly B's own property — which puts B, rather than A, on the hook for costs incurred by the exercise of that agency.

Importantly for our purposes, Viehoff thinks this general phenomenon is at the heart of official immunity. He argues that the compensatory immunity enjoyed by state officials is grounded in the very same considerations that give compensatory immunity to medical decision-making surrogates who loan out their agency to a patient. State officials enjoy compensatory immunity when and because they lend their agency to others.

> **Agency on Loan Account**. A state official is immune from owing compensation for harms imposed as a product of their lending their agency to others. A state official succeeds in lending their agency to someone just in case they satisfy the Deliberative and Justificatory conditions with respect to that person.

Agency on Loan has its virtues. For one thing, it cleanly distinguishes between cases like Fire and Fire 2. In Fire, Firefighter is immune from compensatory liability because the water damage she causes to Neighbor's home is a product of her lending her agency out to Citizen. In Fire 2, by contrast, Citizen has not lent her agency out to anyone. She is acting on her own behalf when she damages Neighbor's home.

A second virtue of Agency on Loan is that it avoids implausible results in a case like Three Fishermen. Bob benefits from Charlie's loss. The

some particular level of reliability (344).

Condition (ii) only requires that B can be expected to benefit by the lending of A's agency given that A deploys that agency with at least that level of reliability it is reasonable for B to expect of A.

simplest versions of the "beneficiary pays" idea predict that Bob thus owes compensation to Charlie. Viehoff's account avoids this result, however, because the damage to Charlie's boat was not the byproduct of anyone's lending their agency out to Bob. This is true even if we imagine a version of the case where Al burns Charlie's boat, not out of jealousy, but purely in order to benefit Bob. Al would still fail to meet the Deliberative Condition, since Al's choices would fail to be guided by the interests of third-parties (in this case, Charlie) to the extent that it would be reasonable for Bob, himself, to be guided by those interests.

Despite these virtues, Agency on Loan suffers serious shortcomings. First, the account is, in some respects, implausibly demanding. In particular, the Deliberative Condition doesn't seem necessary for official immunity. Imagine, for example, a version of Mistaken Imprisonment in which Judge performs his legal duties without treating his agency as a resource assigned to Victim. Imagine that Judge deliberates only on the basis of what the law requires without any mind to Victim's interests. Imagine, moreover, that Judge's concern for what the law requires is conditional on his concern for his own well-being. Judge does what the law requires only because he wants to avoid legal repercussions and because doing as the law requires is the safest path to a nice pension. Judge would thereby fail to satisfy the Deliberative Condition. But (at least for my part) this alteration to the story doesn't budge the intuition that Judge should not be personally on the hook to compensate Victim after his release from prison.

Likewise, consider a variant of Feinberg's oft-cited Cabin case,[18] except where the main character is a state official acting in an official capacity:

> *Fire Cabin*. Jumper is a "smokejumper" with the Forest Service, who has parachuted onto a mountain to prevent the spread of

[18]Joel Feinberg, "Voluntary Euthanasia and the Inalienable Right to Life," *Philosophy and Public Affairs* 7 (1978): 93-123.

> a dangerous forest fire. She is injured in the drop, separated from her team, and can survive only by breaking down the door of Owner's unoccupied cabin to find medical supplies to bandage her wounds. Jumper breaks down the door.

Although Jumper acts without fault in this case, compensation is plainly owed for the damages to the cabin. And yet, just as it is clear that Firefighter should not have to pay compensation for the damages she causes Neighbor in Fire, it is likewise clear that Jumper should not have to pay compensation for the damages she causes Owner in this case. Agency on Loan does not deliver this verdict, however. And that's because Jumper does not satisfy the Deliberative Condition. Her choices to break into the cabin are guided only by facts about her own interests; she isn't lending her agency out to anyone else when it comes to the matter of breaking into the cabin. Agency on Loan, at best, fails to capture the full scope of official immunity.

My second concern is that the Deliberative and Justificatory conditions also fail to be jointly sufficient for compensatory immunity. Consider:

> *Very Neighborly*. Alice spots an ad for an incredible deal on a classic car. She knows her neighbor collects classic cars, and she knows he's been looking for a car like this one. She reasonably judges that Neighbor would buy this car for himself if he were in town, and that it would be in Neighbor's interest for her to buy the car on his behalf. There isn't however, any way to ask for Neighbor's consent that she act for him in this way (Neighbor is out of town and off the grid). She decides to buy the car, but can deliver it to Neighbor's home only by causing $1,000 in damages to another neighbor's lawn.

Imagine Neighbor coming home to find a car in his driveway with a note

attached demanding that he pay the $1,000 in damages. That would be an outrageous demand, and Neighbor would be well within his rights to refuse to accept the car and to refuse to pay for the damages. This is true even if we assume that, had he been home, Neighbor would have bought the car himself even if this meant paying an extra $1,000 to repair someone's lawn.

Viehoff's account, however, appears to predict that Neighbor *is* on the hook to pay for the damages caused by Alice. After all, Alice satisfies the Deliberative Condition: her choices with respect to the car are guided by Neighbor's interests. And she likewise satisfies the Justificatory Condition: Neighbor was not in a position to give his actual consent to the purchase and it would have been in Neighbor's own permissibly pursuable interests to buy the car for himself (both in expectation and *ex post*).[19]

The deeper worry here is that Viehoff's account makes it too easy for others to force us to own things that we may not want to own (precisely because of the liabilities that come with such ownership). I think most people would find it very implausible that Neighbor *owns* the classic car that he finds in his driveway before he performs some act of acceptance. Rather, it seems he *never* owned the car if he refuses to take possession of it when he comes home. But for the same reasons it seems he never owned (even temporarily) Alice's body or agency.

A related worry concerns cases where the beneficiary *is* in a position to refuse actual consent. Imagine a variant of Very Neighborly where Alice is able to get Neighbor on the phone before purchasing the car. She asks Neighbor if she might buy the car for him, but he refuses to borrow her agency in this way. If Alice goes ahead and buys the car for him anyways,

[19]It seems Viehoff's account also delivers the wrong results in the Bells case. I introduce the new case of Very Neighborly above only because it's clearer that Alice satisfies both the Deliberative and Justificatory conditions than that the bell ringer in Bells satisfies both of those conditions.

it doesn't matter that doing so would be in his interest; it would be even more obvious than in the original case that Neighbor has no duty to cover the $1,000 in damages caused by Alice. Refusing to borrow something from someone — whether it be their goat, their car, or their agency — plainly prevents the lending attempt from succeeding. We have very robust powers to prevent others from lending their resources to us.

This fact makes trouble for Agency on Loan because, when we think about official immunity, it just doesn't seem citizens can reliably render state officials liable simply by refusing their aid. If a firefighter attempts to rescue me from a burning building despite my asking to be left for dead, it's very plausible that I prevent her from forcing me to borrow her agency. But it seems very implausible that I thereby succeed in making her personally liable for damages she causes to Neighbor's garden in the process of saving me (certainly not if she is simply following standard department policy by choosing to save me). Given Agency on Loan we should expect our powers to undermine official immunity to be as robust our powers to prevent others from lending us their resources. But this doesn't seem to be the case.

5 The Price of Duty Account

I don't take myself to have delivered anything approaching decisive objections to the various explanations of official immunity canvassed above. That would require a much more extended treatment. I do hope to have shown, however, that there are worries for these accounts that are serious enough to give us good reason to look elsewhere. This section delivers such a look elsewhere.

There is something very appealing about the "beneficiary pays" idea, as considered in §3 and §4. The appeal, I think, is that it is a natural way of precisifying the intuitive idea that official immunity is grounded in considerations of distributive fairness. It offers a way of making sense of the thought that citizens have duties to assume some official liability

because it would be *unfair* if officials had to bear all those burdens on their own. This section proposes an alternative way of unpacking this attractive thought.

I want to start by returning to the earlier quote from Rawls. In that quote Rawls proposes a restricted version of the Beneficiary Pays idea. He doesn't claim that mere benefit puts one under a duty to assume certain costs; he claims only that the benefits of *cooperative projects* put one under a duty of assumption. I argued above that this isn't quite right. There are cooperative projects — like the tolling of the bells — for which beneficiaries are not required to offset costs that fall on the participants. Cost sharing is not required of just any beneficial cooperative project.

That said, I think Rawls's focus on cooperative projects puts us on the right track. There is something special about a certain class of cooperative projects. Consider:

> *Group Rescue*. A hiker's life is threatened by an unexpected snowstorm. Alice and a group of other strangers stumble across him. They can rescue him only by working together to carry him into Owner's cabin. Alice runs ahead to open the cabin, only to find the door and windows locked. She breaks down the door, and the group arrives shortly thereafter to carry the hiker into the cabin. Unfortunately, the hiker's health does not improve and he passes away in the night.

Compare this case with Bells. We noted earlier that Bell Ringer is not owed indemnification by the beneficiaries of the bells. But it also seems that he is not owed any indemnification by the other participants of that project, absent prior agreement to share such costs. Things are plainly different in the case of Group Rescue. The cabin owner is owed compensation for his broken door. But for reasons that are not immediately apparent, it seems

clear that Alice should not have to bear that burden alone. The other members of the rescue team must share this burden with her, regardless of whether there was prior agreement to do so.[20]

This is a surprising difference. The fact that Alice, but not Bell Ringer, enjoys a right to indemnification has nothing to do with the fact that she imposes the harm *permissibly*. Bell Ringer also acts permissibly. Both enjoy a lesser-evil justification for their action — Bell Ringer to save his own life, Alice to save someone else's. Nor can we explain Alice's right to indemnification by appeal to the *receipt of benefits*. Alice's burden must be offset by her co-rescuers, but her co-rescuers benefit neither from Alice's actions nor from the rescue as a whole.

Perhaps it might seem that the key difference is a sort of *complicity*. Although Alice is the person most responsible for the door's destruction, her accomplices are (permissibly) complicit in this trespass by way of their contributing to other aspects of the project of which this trespass is a part. It might seem that they must share in the compensatory burden for the same reason that the getaway driver at a bank robbery might be required to share in compensating those injured in the robbery.[21]

While it is plausible that complicity is one way by which persons can incur duties of indemnification, I don't think this can explain why Alice should not have to bear the compensatory burden alone. For one thing, it's not clear why Alice's co-rescuers would be complicit in the relevant sense in Group Rescue, but Bell Ringer's co-ringers not likewise complicit in Bells. It's not clear that complicity gets us the difference we need between these two cases. Here's a second problem for the complicity proposal. Consider a variant of Group Rescue in which one member of the group that stumbles upon the hiker decides not to help. Call him

[20] At least this cost should be shared on the assumption that other costs associated with the rescue have been distributed fairly.

[21] For a discussion of the concept of *complicitous liability* see Saba Bazargan-Forward, "Complicitous Liability in War," *Philosophical Studies* 165 (2013): 177-195.

Egoist. Egoist's help, let's suppose, is important to the rescue effort, and he could provide it at minimal cost to himself. He is thus under a duty to participate in the rescue. But he would prefer to get home to hot cocoa and a warm fire, and so he leaves the group behind.

Egoist is clearly not complicit in Alice's trespass, since he doesn't participate in the rescue. But just as Egoist has a duty to participate in the rescue in the first place, he likewise seems to have a duty to share in the *costs* of the rescue — regardless of whether he discharges his antecedent duty of participation. Moreover, it seems that his duty of participation is the only thing that can explain his duty of cost sharing: Egoist is morally required to share in the costs of the rescue precisely *because* he was morally required to participate in that rescue in the first place.

The operative principle here, I claim, is what I will call

> **The Price of Duty**. Ceteris peribus, when the Xs are morally required to participate in a cooperative project, the Xs are required to share in certain costs of that project (regardless of whether they actually participate).[22]

This principle of distributive justice is not limited to the *compensatory* costs of morally-obligatory projects. It extends to all sorts of project costs. Suppose, for example, that you and I belong to a small village in a remote desert. A drought has fallen on the region, and all the members of our village will die if they do not get fresh water, and soon. They can get fresh water only if I use my unique construction knowledge to build a small dam and you use your unique chemistry knowledge to build a

[22]Suppose the reason Hiker was in need of rescue was that Villain had injured Hiker by intentionally pushing him off a cliff. Given Villain's culpability for the need for rescue in the first place, it is Villain who should bear the costs of the rescue, and not the rescuers. Cases like these are why Price of Duty is a 'ceteris peribus' principle. More on the qualified status of this principle, and more on the significance of the 'certain' qualifier, in §4.5.

water-purification device. I can build the dam with either mud bricks or concrete blocks. Building with the mud bricks will cost me five days of labor; building with the heavier concrete blocks will cost me ten days of labor. If I use the mud bricks, however, it will make the water much dirtier, and will require greater work on your part. If I use mud bricks, it will take you fifteen days to build your water-purification machine; if I use concrete blocks, it will take you ten days to build your water-purification machine.

Let's assume that you are duty-bound to build the water-purification device, that I am duty-bound to build the dam, and that our duties are equally demanding. Granted this assumption, I find it intuitively compelling that *I owe it to you to build with concrete rather than mud*. If I build the dam with mud, I would wrong you in that I would manage my duty to contribute to the project of providing fresh water for our village in a way that makes it more costly (both comparatively and non-comparatively) for you to discharge your same duty.

In the same way, I claim, the other participants in Group Rescue would wrong Alice if they failed to share in her compensatory burden. By failing to share in that burden, they would manage their duties to contribute to the rescue of Hiker in a way that made it much more costly for Alice to discharge her own duty to do the same. That, I think, is the underlying rationale for the Price of Duty principle. When you and I are morally required to participate in a cooperative project, I am thereby required to bear a fair share of the costs of the project precisely because by failing to do so I would, by my moral non-compliance, make the cost of your moral compliance even more burdensome. You have a right against such treatment.

The Price of Duty principle nicely accommodates our judgments about how the costs of morally obligatory projects should be distributed. It also makes sense of the intuitive difference between cases like Group

Rescue, on the one hand, and cases like Bells or Very Neighborly on the other. In Group Rescue, the damages caused by Alice are the fallout of a project that other people are required to participate in. In Bells and Very Neighborly, by contrast, the damages caused are the fallouts only of projects that are morally *optional* for their participants. There's no moral duty to install and ring bells for your neighborhood to enjoy. There's no moral duty to purchase a classic car for your neighbor.

I find Price of Duty a compelling principle of distributive justice. And it is a principle that nicely explains official immunity. This is because the project of maintaining a just political order is more like the project of rescuing the stranded hiker than like the project of ringing the bells: the citizens of (reasonably just) states are morally required to participate in, and support, this project. Given Price of Duty, it follows that citizens of a state are required to share in certain costs of that project. This demands a certain amount of official immunity, since state officials would bear more than their fair share of the costs of the political project if they were made to bear the full brunt of all the compensatory costs they incur. Not to be confused with the Price of Duty principle itself, let's call this explanation of official immunity that involves an application of that principle the **Price of Duty Account**.

Notice that although this account does not appeal to a version of the Beneficiary Pays idea, benefit still has some role in the story. For one thing, expected benefits can affect whether a project is morally required. The rescuers in Group Rescue, for example, are morally required to participate in the rescue of the hiker only because of this project's prospects for saving the hiker's life. Had the expected benefits been much less, the project would have been morally optional. Likewise, citizens are only required to participate in the political project when and because that project promises to protect certain rights and confer certain benefits (e.g., protection of body and property, a system of corrective justice, access to essential services). Benefits matter in other ways as well. For example,

if someone benefits more from the political project than others (or less), this may affect what share of the costs of that project they must bear. If Alice were the only member of the rescue team in Group Rescue to receive a cash reward for the rescue, this would plainly alter the share of the compensatory costs she should bear.

The fact that state officials benefit the citizenry is thus not irrelevant to official immunity. But benefit plays only an *indirect* role. What fundamentally matters is the fact that the political project is one that citizens are morally required to participate in. Official immunity is just what it looks like to fairly distribute the costs of this mandatory project.

6 Limits

We have some pre-theoretic intuitions about the limits of official immunity. I've relied on some of these intuitions throughout this paper. For example, I've appealed to the widely held belief that a fair and conscientious judge who makes a reasonable mistake in sentencing an innocent person to prison should not have to bear the costs of compensating that victim. This is a clear case of official immunity. And of course there are clear cases of official *non*-immunity. If a police officer takes a baseball bat to the legs of his romantic rival, he is plainly *not* immune from owing his victim compensation, regardless of whether this trespass occurred while he was "on duty".

But between the clear cases of official immunity and official non-immunity is a large penumbra. Direct reflection on cases of official trespass can only get us so far in discerning the *limits* of official immunity. A central contention of this paper is that we can get further by better understanding the *grounds* of official immunity. I've thus far argued that a central ground of official immunity is the Price of Duty principle, which claims that persons who are morally required to participate in a cooperative project are required to share certain costs of that project. Because this principle is a very general principle of distributive justice, that applies in non-political

contexts just as in political ones, we can make progress in thinking about the limits of official immunity by reflecting on much simpler examples of morally required projects. Intuitions are more clear in these cases, I find, and easier to interpret.

When we reflect on such cases, for what sorts of wrongful harm do we find that duties of cost sharing apply? One thing we find is that *duties of cost sharing are not limited to instances of permissible trespass*. Consider a variant of Group Rescue: as the group drags the hiker inside the cabin, Alice misperceives the distance between her body and an expensive glass vase on a nearby table, knocking it over. Destroying the vase is not permissible. And yet, provided we imagine that Alice's misperception is not the result of any negligence or lack of care, it seems that Alice should not have to bear this compensatory burden alone. (These cases also show that duties of cost sharing extend to some harms that do not *contribute* to the obligatory project's end.)

There are four sorts of harm, however, for which duties of cost sharing do *not* seem to extend. First, there is what I will call *incidental* harm. Suppose, for example, that a member of the rescue team decides to enjoy a cigarette after helping pull the hiker inside the cabin. The cigarette slips from his hand, falls between the cracks of the porch, and ignites the dried leaves beneath, causing extensive damage to the porch. These damages seem importantly different from the damages Alice causes to the door. Intuitively, the smoker does not have a claim that his co-rescuers help shoulder the costs of the damage to the porch for the reason that these costs cannot plausibly be described as costs "of" the rescue project. It is not reasonable for the smoker to believe that a cigarette break would contribute in any way to the rescue project, nor does his participation in the rescue project make it difficult or costly to avoid smoking on this occasion.[23] The mere fact that damages are imposed *while* someone is

[23]Intuitively, the costs of the cigarette fire are the smoker's to bear alone, whereas the costs of Alice's accidental breaking of the vase should be shared. What explains this

participating in an obligatory project does not suffice to attribute those damages to the project.

This suggests that official immunity does not extend to harms that are "incidental" to the political project. Citizens do not owe anything to police officers who damage private property while driving to the donut shop. The fact that a police officer is "on the clock" does not suffice to render him immune from such damages; citizens are not required to share in the costs of an official's harm when the harm is unrelated to the official's mission (in the sense that it not reasonable for the agent to believe that the harmful action would contribute to the mission) and the official's mission does not put him in a position that makes it very costly to avoid such harm.

A second type of harm to which duties of cost sharing do not apply is *negligent* harm. Suppose that in Group Rescue, Alice could have simply turned the door handle rather than smash down the door, and that the only reason she didn't try turning the door handle was that she was apathetic to the cabin owner's interests, or that she let her own love of door kicking swamp her concern for the costs this would impose on others, or that she was just being lazy. When we tell the story in this way, Alice clearly lacks a claim against her co-rescuers to their compensatory help. These costs are "on her"; they are costs that were not merely avoidable, but that would not exist had Alice treated others with the care she owed them. The members of an obligatory project do not have a duty to share in the costs that result from another member's failure to act with due care. This makes sense: you cannot reasonably demand that I not make any mistakes in participating in our shared project, but you can reasonably

difference? Two things, I think. First, whereas the choices and actions that most directly lead to the cigarette fire are not plausibly thought of *constituents* of the rescue project, the choices and actions that most directly lead to the broken vase are. Second, whereas the smoker's participation in the rescue project does not put him in a position where he cannot easily avoid creating a fire, Alice's participation in the rescue project does put her in a position where she cannot easily avoid breaking the vase.

demand that I exercise due care to avoid mistakes and that I exercise my agency with a certain level of reliability.

This suggests that official immunity does not extend to negligent harm. Citizens, for example, do not owe anything to the judge who sentences an innocent person to prison because, in his haste to get home for dinner, he looks only at evidence presented by the prosecution and ignores evidence presented by the defense.

A third type of harm for which duties of cost sharing seem not to apply is what I will call *renegade* harm. Suppose Alice and her co-rescuers agree on a plan as to how they will break into the cabin. They all agree on a plan to break a window rather than the expensive door, and they send Alice ahead to do the deed. When Alice arrives at the cabin, however, she decides to deviate from the plan, based on her own belief that it will be easier for the group to use the door than the window. She kicks in the door. This might not be an example of negligence on her part: we might suppose that she breaks down the door precisely because she thinks this will be in everyone's best interest, and we might suppose that she is right about this. Likewise, Alice's action is not disassociated from the project. Her action is connected up with the project in the most straightforward way: it contributes to the project's goal. And yet it seems that Alice cannot demand that her co-rescuers help her bear the extra costs that she imposes by breaking the door rather than the window. This is precisely because she deviated from the plan that was agreed upon by the group.[24]

This suggests that official immunity does not extend to renegade harm. I lack the space to develop a full account of what constitutes renegade conduct by a state official, but the following seems plausible: in a state

[24]It is too strong to say that duties of cost sharing apply only to trespasses that are strictly "part of the plan". These duties also apply to damages that result from *reasonable adaptations* of a plan. For example, Alice would be owed compensatory help if she arrived at the cabin to find the window impenetrable, and for this reason kicked in the door instead.

that has settled legal rules that are the products of democratic procedures, these rules constitute part of the "plan" that state officials must subject themselves to.[25] For example, in the United States it is part of the national plan that federal agents may invade a person's home only if they have a warrant. If a federal agent invades your neighbor's home without a warrant, and his invasion requires him to damage your door, it doesn't matter how careful and efficient the agent acted in the course of the invasion. He owes compensation for the damage to your home, and he is *not* owed indemnification for those costs. The federal agent is not owed indemnification because his trespass is forbidden by the agreed-upon plan for how the U.S. political project is to be pursued. The costs of renegade activity fall outside the scope of project costs that must be shared.

I'll conclude this section by considering an especially challenging case: the enforcement of unjust laws. These are not cases in which an official deviates from the plan. These are cases in which an official follows the plan, but where the plan itself is mistaken. Must the citizenry share in the costs of the enforcement of unjust laws?

Reflection on simple cases of group rescue suggests that some, but not all, of the costs of the enforcement of unjust laws should be shared by the citizenry. Suppose nine of us are planning how to rescue the hiker. Three plans are proposed. Four members propose that we break into the unoccupied cabin. Four other members propose that we attempt to bring hiker down the mountain on a sled. Neither of these two options is obviously better than the other, in light of the common information available to the group. The ninth member of the group proposes a surprising third plan. He proposes that we throw the hiker off a cliff, in his sincere belief that angels will catch the hiker and fly him to safety.

[25] See Scott Shapiro, *Legality* (Harvard, Massachusetts: Harvard University Press, 2011) for a view according to which a system of law *just is* a sort of plan.

Suppose that the ninth member of the group — a true master of the rhetorical arts — somehow manages to convince a majority of the group to go along with his plan. Eight members of the group proceed to throw the hiker off the cliff, with only one person refusing to participate. As it turns out, the hiker is not caught by angels, but instead falls to his death, his body crashing into and injuring a bystander at the bottom of the cliff. This much seems clear: compensation is owed for the bystander's injuries, but the one person who did not participate in the killing is *not* required to share in this compensatory burden.

But now suppose instead that the group decides on the plan to break into the cabin. Eight members of the group go ahead and break into the cabin, which requires breaking down the cabin door. One person abstains from participating. As it turns out, however, the majority implements a suboptimal plan: had they chosen to take hiker down the mountain on a sled, they could have rescued the hiker without any trespass whatsoever. Does this mean, as in the previous iteration of the case, that the lone non-participant needn't share in the compensatory burden? Here things seem very different. It seems the dissenter *should* share in the compensatory burden, precisely because this burden is the result of a reasonable mistake. Given the information publicly available to the group, it was reasonable for them to disagree as to whether the cabin-invasion plan or the sledding plan was best. By contrast, it was not reasonable for them to disagree as to whether it was best to toss the hiker off the cliff.

This is suggestive of a way to think about official immunity in the context of the enforcement of unjust laws. Official immunity extends to the costs of the enforcement of *reasonable*, unjust laws. Some unjust laws are "reasonable" in the sense that, although they are in fact unjust, it is reasonable for people to disagree about whether they are just, relative to some body of public information.[26] Citizens have a duty to share in the compen-

[26]Here I won't take a stand on the contentious question what this relevant body of public information is, and on what constitutes a reasonable position relative to that body

satory costs of the (non-negligent) enforcement of such laws, regardless of whether they are among those who believe these laws to be just. But where it is not reasonable to defend an unjust law, citizens will not have a duty to share in the compensatory costs of the enforcement of such laws. State officials do not enjoy immunity for instances of *unreasonable* harm.

Some seriously unjust laws are subject to reasonable disagreement; some mildly unjust laws are not. Because the degree of reasonability and the degree of injustice come apart in this way, the above considerations imply that we can't infer whether an official is immune from the enforcement costs of some law on the basis of the degree of the injustice of the law. Likewise, the above considerations imply that state officials are not guaranteed immunity simply by the non-negligent execution of their role. Citizens do not owe it to state officials to share in the compensatory costs of their enforcement of unreasonable laws, even if those officials are not in any way to blame for their actions.

7 Conclusion

The Price of Duty Account tells us that state officials enjoy a sort of vicarious immunity with respect to many sorts of harm. Many of the costs imposed by state officials should be shared by the citizens who are morally required to participate in the political project of which those costs are a fallout. Those citizens have a moral duty to indemnify state officials for those costs. But this isn't true for *all* of the costs imposed by state officials. There are exceptions. We noted four. Official immunity does not extend to instances of incidental, negligent, renegade, or unreasonable harm. That the Price of Duty Account suggests these exceptions is not just a feature, but a virtue, of the account. These are plausible exceptions. The account fills out the penumbra in our pre-theoretic judgments about the limits of official immunity in a way that does not conflict with our more determinate judgments about official immunity.

---

of information.

On this account, the *moral* immunity enjoyed by state officials looks quite different from the *legal* immunity they enjoy in most contemporary nations. In my own country of the United States, for instance, legal immunity outruns moral immunity along a number of fronts. U.S. law grants "absolute immunity" to a number of offices. Judges, prosecutors, legislatures, and certain members of the executive branch are almost never made to pay compensation, even if they've acted negligently or unlawfully.[27] Most other officials enjoy "qualified immunity", under which officials can be made to pay compensation for the wrongful harms they cause, but only if they have acted in a way that any "reasonable person" would know violated "clearly-established" law.[28] Even this standard, however, extends legal immunity well beyond the bounds of moral immunity, as conduct can easily impose incidental or negligent damages without passing the 'clearly-established' and 'reasonable person' tests.[29] And even where the law does hold an official liable to pay compensation, many such officials enjoy indemnification. This is especially true of law enforcement officers, who, even when they are not protected by qualified immunity, almost always have their compensatory burden assumed by their department or local/state government.[30] This, of course, is just an indirect way of shifting an officer's compensatory burden onto the taxpayers.

The fact that legal immunity in the United States outruns moral immunity is not yet to say that U.S. immunity policy is deficient. But it is to say that

[27]Fred Smith, "Local Sovereign Immunity," *Columbia Law Review* 116 (2016): 411.

[28]*Harlow v. Fitzgerald*, 457 U.S. 800, 818 (1982).

[29]Courts have generally interpreted the "clearly-established" test as a *very* demanding one. "The Court has said that, except in extraordinary circumstances, the law is clearly established only if a prior case has declared the conduct unconstitutional. And that prior case must have facts that map neatly onto the facts of the plaintiff's case" (Joanna Schwartz, "Qualified Immunity's Boldest Lie," Chicago Law Review 88 no. 3 (2021): 607). See also Tyler Finn, "Qualified Immunity Formalism: 'Clearly Established Law' and the Right to Record Police Activity," Columbia Law Review, 119 no. 2 (2019): 445-486; and John Jeffries, Jr., "The Liability Rule for Constitutional Torts," Virginia Law Review, 99 no. 2 (2013): 207-270.

[30]Joanna "Schwartz, Police Indemnification," NYU Law Review 89 no. 3 (2014).

there is a burden to justify this divergence. When the state grants legal immunity for an official's incidental, negligent, renegade, or unreasonable harm, the state chooses not to enforce a moral duty of compensation that is owed by that official. This is pro tanto wrongful: the state has a pro tanto duty to ensure the satisfaction of the compensatory claims of its citizens, and this duty is all the more demanding in the case of official harm, as the state typically bears some complicity for the harm done by its own agents. The upshot: the state had better have a compelling justification for the ways in which it extends legal immunity beyond the limits of moral immunity.

It is a complex and largely-empirical question whether there is such a justification for this divergence. I think there is a prima facie case for skepticism, however. The standard arguments in favor of such extensive immunity regimes as we have in the United States purport that a less extensive regime would deter officials from doing their jobs well, deter quality candidates from applying for state offices, and tie up essential government workers and resources in court.[31] But there is a growing consensus amongst legal scholars that these arguments are empirically untenable, and that official immunity could be scaled back in the United States without serious overdeterrence or administrative costs.[32] Perhaps there are better arguments out there. But notice the hurdle these arguments need to clear. It wouldn't be enough to show that the benefits of a more expansive immunity policy are greater than the benefits of a less extensive one (a challenge in its own right). It must also be shown that this difference in benefit is so great as to outweigh the state's pro tanto

[31]See the SCOTUS opinion in *Harlow v. Fitzgerald*, 457 U.S. 800, 814 (1982).

[32]This is largely due to the comprehensive work of Joanna Schwartz. In addition to her "Indemnification" and "Boldest Lie", see also "How Qualified Immunity Fails," *Yale Law Journal* 127 no. 2 (2017): 2-76; "The Case Against Qualified Immunity," *Notre Dame Law Review* 93 no. 5 (2018): 1797-1851; and "After Qualified Immunity," *Columbia Law Review* 120 no. 2 (2020): 309-383. For a partial defense of qualified immunity see Aaron Nielson and Christopher Walker, "A Qualified Defense of Qualified Immunity," *Notre Dame Law Review* 93 no. 5 (2018): 1853-1885.

duty to enforce the moral duties of its own officials.