

Original citation:

Tadros, Victor. (2018) Causal contributions and liability. *Ethics*, 128 (2). pp. 402-431.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/89365>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

© 2018 by the University of Chicago. <https://doi.org/10.1086/694275>

A note on versions:

The version presented in WRAP is the published version or, version of record, and may be cited as it appears here.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

Causal Contributions and Liability*

Victor Tadros

This article explores the extent to which the magnitude of harm that a person is liable to suffer to avert a threat depends on the magnitude of her causal contribution to the threat. Several different versions of this view are considered. The conclusions are mostly skeptical—facts that may determine how large of a causal contribution a person makes to a threat are not morally significant, or not sufficiently significant to make an important difference to liability. However, understanding ways in which causation may be scalar helps to deepen our understanding of other morally significant facts, such as responsibility.

Many just war theorists believe that at least one ground of a person's liability to be harmed to avert an unjust threat is her moral responsibility for causing the threat. If a person is liable to be harmed to avert a threat, that person lacks a right not to be harmed to avert that threat. Harming that person to avert the threat does not wrong that person, and this makes a significant difference to the permissibility of harming that person. Thus, the view under consideration implies that whether a person causes an unjust threat determines whether it is permissible to harm that person.

Suppose this is true. Does the magnitude of a person's causal contribution to a threat make a difference to liability? If causation grounds liability and causal contributions are scalar, it is initially plausible that liability depends on the magnitude of the causal contribution a person makes to a threat. Some just war theorists rely on this idea to support a reasonably general, if not uniform, principle of noncombatant immunity.

This issue has great practical importance. Causal contributions seem to vary widely across combatants and noncombatants. Combatants in the

* Earlier versions of this article were presented at Stockholm University, Oxford University, and a workshop on causation, responsibility, and war in Catania. I am grateful to the participants at those events, and especially to Tom Dougherty, Cécile Fabre, Helen Frowe, Alex Kaiserman, Carolina Sartorio, and Adam Slavny, as well as two referees and the associate editors of *Ethics*. I am grateful to the Leverhulme Trust for a Major Research Fellowship that supported the writing of this article.

same unit as the combatants who directly pose threats support their missions with navigational and tactical advice, they encourage combatants to pose these threats, they help to load weapons and identify targets, and so on. Combatants beyond the unit also contribute—without them the war would not have occurred, they secure the territory from which further advances can be made, they divert threats away from the unit, they provide intelligence, they provide a disciplinary hierarchy which makes the war machine more effective, and so on. Pilots, drivers, medics, caterers, and administrators contribute to the infrastructure of war, both at home and abroad. And wars are conducted only because of the large-scale civilian effort that supports them: manufacturing munitions and vehicles, military training, building roads, developing weapons, sending food and clothing, gathering and assessing intelligence, providing moral support, voting for the war, governing it, and so on, are all essential to the war effort.

To assess whether causal contribution matters, and if so, how, we need a clearer understanding of how causal contributions are scalar, if they are. This will allow us to judge whether any facts that make them scalar are morally significant. As there are many causal theories, there are many possibilities.

My main aim is not to determine whether, or how, causal contributions are scalar. It is to assess the implications of different ideas about how causal contributions might be scalar for liability to defensive harm. My conclusions are mostly skeptical. In most cases, facts that might make causal contribution scalar are morally irrelevant. In some cases where they are relevant, the facts that might make a causal contribution scalar are already well recognized as relevant to liability. In that case, the view that liability depends on the magnitude of a person's causal contribution may deepen our understanding of familiar ideas, but it does not provide a new character on the moral landscape. There are some cases where my conclusions are more complex. None of the ideas I consider significantly supports a reasonably general principle of noncombatant immunity. Of course, that does not show that there is no such principle, only that it is not grounded in differences in causal contribution.

Section I outlines the role that causal contribution has played in just war theory. Section II makes some clarificatory and methodological remarks. Section III considers the idea that a person might contribute part of a threat rather than a whole threat. Section IV focuses on the idea that a person might contribute part of the cause rather than the whole cause. Section V is concerned with the idea that the magnitude of a person's causal contribution to a threat depends on the relationship between the threat and features of personhood. Section VI considers the distinction between enabling and direct causation. Section VII explores David Lewis's counterfactual account of degrees of causal contribution. Section VIII considers overdetermination and preemption cases. Section IX evaluates prob-

abilistic theories of causation. Section X explores proximity and remoteness.

I. CAUSAL CONTRIBUTIONS IN JUST WAR THEORY

Several just war theorists, such as Cécile Fabre, Jeff McMahan, and Seth Lazar, rely on the idea that liability to be harmed depends on the magnitude of causal contribution to support a general, if not uniform, principle of noncombatant immunity.

Fabre argues that a person who contributes to a wrongful venture resulting in unjust lethal threats is liable to be killed to avert those threats only if her causal contribution meets some threshold. She writes,

A contribution must, on its own individual terms, meet a threshold of causal significance in order for its author to be liable. Tightening screws on tank engines, testing the sweat-absorbing capacities of the clothes which soldiers will wear in the desert, and adjusting the speed level of food packaging machines do not, it seems to me, pass the threshold. Nor, for that matter, does designing a tiny piece of equipment which goes into a gun do so. By contrast, taking overall responsibility for negotiating and drafting sales contracts between one's factory and the army might; so might driving a truckload of munitions or protective clothing to an armory division, and so on.¹

She does think that there is an exception, which we will consider below, where a person causally contributes a small portion of a larger threat. In that case, she thinks that the person's liability depends on the size of the whole threat, not the size of the part that the person contributes. So she would now qualify, but not reject, the first sentence in the quote above.²

McMahan argues that in practice the responsibility of civilians for unjust threats is typically too trivial to make them liable to be killed. He writes,

Most civilians have, on their own, no capacity at all to affect the action of their government. They may pay their taxes, vote or even campaign for particular political candidates (sometimes on the basis of general sympathy with their overall positions on matters of policy but seldom because of their advocacy of war), participate in the culture from which the country's political leaders have emerged, fail to protest their country's unjust war, perhaps because they correctly be-

1. Cécile Fabre, "Guns, Food, and Liability to Attack in War," *Ethics* 120 (2009): 36–63, 61. See also Cécile Fabre, *Cosmopolitan War* (Oxford: Oxford University Press, 2012), 76–77, and *Cosmopolitan Peace* (Oxford: Oxford University Press, 2016), 6.

2. Fabre, *Cosmopolitan War*, 76–77.

lieve that to do so would be ineffective, or perhaps because they approve of the war, and so on; but none of these things, nor even all of them together, is ordinarily sufficient for the forfeiture of a person's right not to be attacked and killed. Military attack exceeds what a person may ordinarily be liable to on the basis of these comparatively trivial sources of responsibility.³

McMahan does not make the idea that these sources of responsibility are comparatively trivial completely transparent. He might seem to refer to the mental states of the civilians—their intentions, knowledge, or beliefs. But this reading is not best. He notes that some of these people intend their contributions, and others knowingly contribute. Furthermore, he thinks that even those who are minimally responsible for causing or posing lethal threats to others are liable to be killed to avert these threats.⁴ A better reading is that McMahan believes that the causal contributions that these civilians make to lethal unjust threats are too small to make them liable to be killed.⁵ Later, McMahan relies on the same idea as part of a package of considerations that make it wrong to attack civilians in a wide range of circumstances in practice.⁶

Neither Fabre nor McMahan provides a clear account of how causation is scalar, or why this matters to liability. Fabre, for example, simply suggests that certain functions, such as tightening the screws on tanks, make a smaller causal contribution to a threat than others, such as driving a truck-load of munitions. And legal scholars also take this approach.⁷ McMahan rejects the significance of proximity.⁸ His first sentence in the above quote might suggest that he believes that if X cannot on his own affect whether Y poses a threat, any causal contribution that X makes to the threat that Y poses is too small for X to be liable to be killed. But he does not explain why.

Seth Lazar thinks that liability to be killed can be grounded in either high culpability with minimal causal contribution or minimal culpability with great causal contribution. He briefly suggests a set of conditions that

3. Jeff McMahan, *Killing in War* (Oxford: Oxford University Press, 2009), 225.

4. See Jeff McMahan, "The Basis of Moral Liability to Defensive Killing," *Philosophical Issues* 15 (2005): 386–405.

5. For confirmation that this is the right reading, see Jeff McMahan, "Who Is Liable to Be Killed in War," *Analysis* 71 (2011): 544–59, 548–50. McMahan acknowledges that he does not have a clear account of causation in "Basis of Moral Liability," 396.

6. See McMahan, *Killing in War*, 231.

7. The use of such examples is common. See, e.g., Michael N. Schmitt, "Humanitarian Law and Direct Participation in Hostilities by Private Contractors or Civilian Employees" (expert paper submitted to the International Committee of the Red Cross, October 2004, 14–24); Yoram Dinstein, *The Conduct of Hostilities under the Law of International Armed Conflict* (Cambridge: Cambridge University Press, 2016), 178–81.

8. McMahan, "Basis of Moral Liability," 396.

must be fulfilled for the person to make a great causal contribution to a threat: her act must be necessary, sufficient, and direct. But he does not explain what this means or defend his view further.⁹ He does provide one example: he thinks that munitions workers make an insufficiently large causal contribution to be liable to be killed to avert lethal threats in war unless they are culpable. He thinks that it is too difficult to tell whether they are culpable to justify killing them.¹⁰

Helen Frowe criticizes the general view that those who make small causal contributions are not liable to be killed by considering the following example:

Mob: Mafia Boss wants to take Victim out, but he cannot afford to hire Assassin, who is extremely skilled and thus extremely expensive. Mafia Boss has a whip-round amongst all the members of his mob, none of whom really like Victim. Everyone coughs up a few pounds for the assassination fund.¹¹

Frowe believes that each contributor is liable to defensive harm even though his causal contribution to the hit is small. But even if Frowe's intuition about liability is sound, it is not clear why she thinks each person's contribution is small. Overall, then, just war theorists have relied on different views about the magnitude of causal contribution without any deep investigation of causation.

II. CONTEXT AND METHODOLOGY

Here are some clarifications and methodological suggestions. First, I work on the controversial but widely accepted assumption that causation more generally grounds liability. This is for the obvious reason that if causation does not ground liability, neither do causal contributions.

Second, we are only concerned with whether causal contributions affect liability. There are related issues, such as whether the stringency of our duties not to contribute to unjust threats depends on the magnitude of our causal contribution to those threats. Although these issues are related, their relationship may not be straightforward, and I do not explore it.

Third, there are different possible views of how causal contributions relate to liability. A threshold view is that liability to be harmed to avert a threat depends on some causal threshold being reached. A scalar view is that the magnitude of harm that a person is liable to suffer to avert a lethal threat depends on the magnitude of that person's causal contribution to the threat. A scalar view is more attractive. This is because on

9. Seth Lazar, *Sparing Civilians* (Oxford: Oxford University Press, 2015), 94–95.

10. *Ibid.*, 95.

11. Helen Frowe, *Defensive Killing* (Oxford: Oxford University Press, 2014), 175.

the threshold view small differences in the magnitude of a person's causal contribution can make a very large difference to the magnitude of harm that a person is liable to suffer to avert a threat near the threshold.

Fourth, as causation may be a complex idea with many dimensions, there may be different kinds of causal contributions that a person can make. Those with this view of causation may think that only some of these contributions are relevant to liability, or that only some come in degrees in a way that affects the magnitude of liability.

So the view that I will consider is the following:

Causal Contribution: The magnitude of harm that a person is liable to suffer to avert a threat because of the causal contribution that she makes to that threat is proportionate to the magnitude of the causal contribution that she makes to the threat, or to some dimension or dimensions of the causal contribution she makes to the threat.

There are too many versions of this view to conduct an exhaustive study, so I consider only those with the greatest initial plausibility.

Fifth, even if causation affects liability, it is only one among a range of factors that determine liability. The other most important factor is responsible agency. At one extreme, there are those who are highly culpable for the threats they contribute to. At the other, there are those who causally contribute to a threat but are not morally responsible at all. And there are grades of responsibility between these extremes.

A simple view is that *Causal Contribution* applies equally across the full spectrum of cases. But some might think that *Causal Contribution* and responsibility interact in a more complex way. One view is that high levels of culpability swamp differences that degrees of causal contribution would otherwise make to liability, but low levels of culpability do not.¹²

Although this is a possible view, it seems doubtful. Those who believe that causation grounds liability typically believe this even for those who are highly culpable: they think that those who are highly culpable for attempting to pose threats are liable to less harm than those who actually pose threats.¹³ Suppose that this is right, that causation is scalar, and that *Causal Contribution* applies to those who are not very culpable.

In light of this, it is hard to believe that great culpability swamps the significance of causal contributions. To see this, suppose that X makes a

12. Thanks to a reviewer of *Ethics* for prompting me to consider this. It is not completely clear, but Seth Lazar might intend something like this view. See Lazar, *Sparing Civilians*, 94.

13. See, e.g., McMahan, "Basis of Moral Liability," 386; Victor Tadros, "Causation, Culpability, and Liability," in *The Ethics of Self-Defense*, ed. Christian Coons and Michael Weber (Oxford: Oxford University Press, 2016), 110–30. Lazar is also sympathetic to the view that causal contribution is necessary for liability to be killed. See Lazar, *Sparing Civilians*, 94.

miniscule contribution to a lethal threat and Y makes no contribution to that threat. The view under consideration implies that if X and Y are minimally culpable, there is little difference in their liability. If they are highly culpable, there is a large difference in their liability. This is initially unattractive. Nevertheless, I will consider both cases of high culpability and those where culpability is less.

I leave aside cases of moral responsibility without culpability and causal contribution without responsibility. A person who is morally responsible for causing an unjust threat is a person who permissibly takes a risk of posing an unjust threat, and that risk is realized. A person who lacks responsibility for posing a threat either has no reason to believe that she will or might pose a threat or is not in control of whether she does so.

Some believe that some or all of these people are liable to be harmed to avert threats they pose. As it is difficult to assess liability in these cases independently of any questions about the magnitude of a person's causal contribution to a threat, we should not be confident in our assessment of the magnitude of a person's causal contributions to a threat in such cases either. As many people who causally contribute to unjust lethal threats posed in war are at least somewhat culpable for doing so, even this more restricted study has wide-ranging implications. With this in mind, let us turn to different versions of *Causal Contribution*.

III. PORTION OF THE THREAT

The idea that causal contributions come in degrees has not received a great deal of attention in the philosophy of causation.¹⁴ Some might be skeptical about this idea in general. One source of skepticism arises from the still-popular counterfactual view of causation. Consider this simplistic counterfactual view: one event, E1, causes another event, E2, if and only if E2 counterfactually depends on E1. Obviously, this is a woeful theory of causation—even if causation ought to be understood counterfactually, counterfactual dependence is neither necessary nor sufficient for causation. But it plausibly explains simple cases.

As counterfactual dependence is not scalar, some might doubt the idea that there are degrees of causal contribution. Either E2 counterfactually depends on E1 or it does not. For example, on the simplistic

14. For some exceptions, see David Lewis, "Causation as Influence," *Journal of Philosophy* 97 (2000): 182–97, 189–91; Matthew Braham and Martin van Hees, "Degrees of Causation," *Erkenntnis* 71 (2009): 323–44; H. L. A. Hart and Tony Honoré, *Causation in the Law*, 2nd ed. (Oxford: Oxford University Press, 1985), 233; and Michael S. Moore, *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics* (Oxford: Oxford University Press, 2009). For critical discussion of Moore, see Helen Beebe, "Legal Responsibility and Scalar Causation," *Jurisprudence* 4 (2013): 102–8.

view, if a threat counterfactually depends on X's act, X causes the threat; if it does not, X does not.

Even on the simplistic view, though, causal contributions can vary. X poses a threat if X will cause harm unless X is prevented from doing so. The threat is the total harm that X will cause. Part of the threat might counterfactually depend on X's conduct. The degree of X's causal contribution might correspond to the size of the part that counterfactually depends on X's conduct.¹⁵ This general idea could then be adjusted in light of more sophisticated counterfactual views. *Portion of the Threat* is a version of *Causal Contribution* that claims that liability depends on the magnitude of the portion of the threat that counterfactually depends on one's act (or some more sophisticated proposal of the relevant counterfactuals).

This case illustrates the view:

Flamethrower: Y's flamethrower is running low on fuel and X tops it up. Y is about to fire at V, threatening to cause V third-degree burns. Had X not acted, all else would have been equal, but Y would have threatened to cause V second-degree burns.

Portion of the Threat implies that X's contribution is the difference between second-degree burns and third-degree burns and his liability depends on that amount. He is thus liable to suffer less harm to avert the threat posed by Y than he would have been had he completely refueled Y's empty flamethrower.

The general view that the magnitude of a person's liability to preventive harm depends on the magnitude of harm that the person threatens to cause is highly plausible. In light of this, some might find it obvious that *Portion of the Threat* is true. But some argue that if many people each make a very small contribution to a very large threat, the liability of each depends on the magnitude of the whole threat, rather than the part that counterfactually depends on their conduct. Fabre, for one, believes roughly this.¹⁶ The issue is too difficult to resolve here, but there is at least some doubt that *Portion of the Threat* is true.

Even if *Portion of the Threat* is true, a person can become liable to be killed by making a small contribution to a large threat: the threat can be to many lives, and so the portion of the threat that the person contributes may be a matter of life and death to one or more people.

15. This sense in which there can be degrees of causal contribution is acknowledged in Stephen Mumford and Rani Lill Anjum, *Getting Causes from Powers* (Oxford: Oxford University Press, 2011), 21.

16. See Fabre, *Cosmopolitan War*, 76–77. For further discussion, see Jeff McMahan, "Proportionality in Defense against Inflictors of Small Harms" (unpublished manuscript).

IV. PORTION OF THE CAUSE

On some views of causal contribution, a threat, T, may counterfactually depend on each of two earlier events, E1 and E2, and yet E1 may make a greater causal contribution to T than E2. One idea is that the causal contribution that one makes to a threat depends on the magnitude of the contribution that one makes to the cause of the threat. One version of this idea is that the magnitude of one's causal contribution corresponds to the magnitude of the portion of the cause that one contributes.

Consider the following:

Blood Loss: X and Y independently cut V, causing him to lose blood. V dies of blood loss. V would not have died had he not been cut by both X and Y. Y's cut causes V to lose much more blood than X's cut.

The whole threat counterfactually depends on both X's act and Y's act. But Y seems to make a greater causal contribution to the death than X.¹⁷ One plausible explanation is that the death is caused by something quantifiable—the loss of sufficient blood—and Y contributes a greater portion of the cause than X. *Portion of the Cause* is the related version of *Causal Contribution* that claims that liability depends on the size of the portion of the cause that one contributes. This theory only applies when the cause of harm can be divided into portions and there is a principled way of determining the size of these portions.

Some might find it intuitive that Y is liable to suffer more harm than X to avert the threat they both cause. This view might be supported by the intuition that if V's death could be prevented by killing either X or Y, Y should be killed. *Portion of the Cause* provides one explanation of any such intuition.

I have at most, though, a weak intuition that Y should be killed. Furthermore, *Portion of the Cause* is not the only possible explanation of the intuition. Here is an alternative. Had Y not acted, X would have caused V less harm than Y would have caused V had X not acted. This fact might be relevant to liability. We can remove this distraction by considering a case where no harm would have been caused by either X or Y had they acted alone, but together their acts contribute to the cause of a great deal of harm:

Boulder: X and Y together push a boulder off a cliff that crushes V. Neither would have been able to get the boulder off the cliff without the other pushing. X is much stronger than Y and exerts much more force on the boulder.

17. See Moore, *Causation and Responsibility*, 275–76.

The force that X and Y together provide causes the boulder to roll off the cliff, killing V. The existence of this magnitude of force counterfactually depends on both X pushing and Y pushing. But X provides a larger portion of the cause than Y. There seems to be no difference in the liability of X and Y, and this seems true whether X and Y are highly culpable or they are merely negligent. Thus, I doubt that *Portion of the Cause* is right.

V. PERSONAL CAUSATION

In some cases, causal contributions seem to depend on the relationship between a person and an effect. When I causally contribute to an effect, the relevant causal relationship is between the effect and me. This relationship can vary in strength, it might be argued, and a person's causal contribution depends on it.

To see that this is plausible, consider a case not involving wrongdoing:

Performance: Y puts on a moving musical performance which X records. Later, the recording is played to an audience, whose members are moved.

The audience being moved counterfactually depends on both X making the recording and Y performing, and both causally contribute to the audience being moved. But Y seems to make a larger causal contribution than X. This is true even if X and Y each intend the audience to be moved. But there is nothing quantifiable that Y provides more of than X.

One explanation is that the objective probability of Y's act resulting in the audience being moved is greater than X's. But even if we can make sense of objective probabilities, this is not true if X is uncertain to record the performance but Y is certain to perform. A second explanation is that X is easily replaceable whereas Y is not. But the intuition holds even if there are many musicians to replace Y and hardly any sound engineers to replace X.

Personal causation provides a more immediate and appealing explanation. It exists in cases where there is the right kind of relationship between properties central to personhood and the effect. Properties such as agency are central to personhood, and that explains why we are powerfully inclined to say that a person caused a result where the person's agency was involved in the right way. This is in contrast to cases where my body is involved in bringing about an effect but I am not, such as cases of involuntary muscle spasm. In such cases, my body causes the effect, but I do not.

The line between me doing something and my body doing it is not sharp; what I do while sleepwalking, for example, is a borderline case. That explains why we are less inclined to say that a person caused an

event where the person was sleepwalking. Furthermore, consciousness and agential involvement come in degrees, and the extent of a person's causal contribution to an effect may depend on these degrees.

Where two people contribute to an event, the magnitude of each person's causal contribution may depend on the depth of the different relationships between features central to personhood and the event. In *Performance* the relationship between features central to Y's personhood and the audience being moved is plausibly deeper than the relationship between features central to X's personhood and the audience being moved. This explains why we are inclined to see Y as making a greater causal contribution to the audience being moved than X.

Different versions of this idea focus on different properties of personhood. And different versions focus on different claims about the causal relationships between properties of personhood and the effect. Furthermore, some versions claim that the same properties are relevant in all cases, whereas others claim that this is a contextual matter that depends, for example, on the nature of the activity we are concerned with. Here I rely only on the rough idea.

Now consider the following:

Torture: X fixes the door to a torture room. Y tortures V in the room. Z stands by and cheers Y on. Had any of X, Y, or Z not done what they did, V would not have been tortured.

Y's causal contribution to V's suffering seems greater than X's or Z's, even though V's suffering counterfactually depends on all of these acts. The deeper connection between the torture and features of Y's personhood is one explanation.

Personhood is a version of *Causal Contribution* that claims that a person's liability depends on the magnitude of her personal causal contribution. To see that this view is attractive, suppose that harming X, Y, or Z could save V and that all three intend the torture. Y, it seems, should be harmed, other things being equal. The fact that Y's agency is more deeply engaged in bringing about the torture attractively explains this result.

Even with only a rough characterization of *Personhood* in hand, we can also see that it might help to explain some intuitions in war, for example, why those who fire their weapons seem liable to a greater degree of harm than those who load them. Firing does seem more deeply connected to features of the agency of the person who fires than to those of the person who loads.

The main contribution that *Personhood* promises to make to the morality of war, though, is to deepen the familiar idea that liability is grounded in moral responsibility. Furthermore, even if *Personhood* is

right, it provides little support to noncombatant immunity. As defenders of *Causal Contribution* are well aware, responsibility-relevant facts do not powerfully distinguish combatants and noncombatants. Lethal threats may be just as deeply attributable features central to the personhood of noncombatants as to the personhood of combatants. Noncombatants may exercise special skill in performing their roles, for example, in the manufacture of precision weapons, in training combatants, and so on. And they, like combatants, often intend to contribute to the war effort.

Furthermore, many combatants are relatively unskilled, and the threats they pose are not deeply connected to central features of personhood, but this does not seem to undermine their liability to be killed. A poorly trained and terrified combatant who starts to fire an automatic weapon into a crowd without thinking about it is surely liable to be killed to avert the threat he poses. Still, *Personhood* may well be both true and important, and it is a version of *Causal Contribution*.

VI. CAUSING AND ENABLING

Enabling is a version of *Causal Contribution* that claims that those who enable threats are liable to less harm than those who directly cause them.¹⁸ The distinction between enabling and causing is intuitive but hard to explicate. Consider the following:

Match: A damp match is dried and then struck.

The striking seems to cause the match to be lit, whereas the drying only enables this.¹⁹ But why?

18. Adil Ahmad Haque relies on the distinction between enabling and causing to help support a more stringent duty not to kill noncombatants; see Adil Ahmad Haque, *Law and Morality at War* (Oxford: Oxford University Press, 2017), 73–77. There is a large literature on whether enabling harm is easier to justify than causing harm. See, e.g., Jeff McMahan, “Killing, Letting Die, and Withdrawing Aid,” *Ethics* 103 (1993): 250–79; Kadri Vihvelin and Terrance Tomkow, “The Dif,” *Journal of Philosophy* 102 (2005): 183–205; Samuel C. Rickless, “The Moral Status of Enabling Harm,” *Pacific Philosophical Quarterly* 92 (2011): 66–86; Christian Barry and Gerhard Øverland, “The Feasible Alternatives Thesis: Kicking Away the Livelihoods of the Global Poor,” *Politics, Philosophy and Economics* 11 (2012): 97–119; Fiona Woollard, *Doing and Allowing Harm* (Oxford: Oxford University Press, 2015), chap. 4; Jason Hanna, “Enabling Harm, Doing Harm, and Undoing One’s Own Behavior,” *Ethics* 126 (2015): 68–90. My own sympathies are with the view that enabling harm is no easier to justify than causing it and that the cases where this seems false are mainly those that involve withdrawing resources that one is entitled to.

19. See Lawrence Brian Lombard, “Causes, Enablers, and the Counterfactual Analysis,” *Philosophical Studies* 59 (1990): 195–211, 201–2. For doubts, see Penelope Mackie, “Causing, Enabling, and Counterfactual Dependence,” *Philosophical Studies* 62 (1991): 325–30. For a reply, see Lawrence Brian Lombard, “Causes and Enablers: A Reply to Mackie,” *Philosophical Studies* 65 (1992): 319–22.

The contrast is not simply one of proximity. Events in a long chain may directly cause each other. The first event, in this case, is a direct cause of the last, as is true of a long chain of dominoes, each of which topples the next. A better view is that enablers are distinct from prior causes in that they provide background conditions that make possible the relationship between a cause and its outcome.²⁰ But it is difficult to explain what makes something a mere background condition, and the difficulty of doing this may suggest that there is no real metaphysical distinction between causing and enabling, but rather a difference in the conversational felicity of pointing out that something is a cause when this does not illuminate the case under consideration.²¹

Without an analysis of the distinction in hand, let us explore whether the intuitive distinction between causing and enabling significantly affects liability. I doubt that it does. First, culpably enabling a lethal threat is sometimes sufficient for liability to be killed. Consider the following:

Burn: X and Y hate V. Y captures V and ties him to a tree to burn him to death, but he can't get his wet match lit. X has a battery-powered hair dryer with him and begins drying the match, which Y will then strike to burn V to death. D can prevent V from being burned to death only by killing X.

X seems liable to be killed to prevent V's death even though X only enables Y to kill him.

Perhaps it might be argued that even if X is liable to be killed, his liability is less than Y's. This view might be supported by noting that if D could prevent V from being killed by killing either X or Y, he ought to kill Y. This view is intuitive, but I doubt that it is best explained by the distinction between enabling and direct causes.

Who ought to be killed does not seem powerfully to depend on who dries and who strikes the match, but rather on other features of the case, such as the fact that the killing is mainly Y's plan and that his personal involvement is greater. If we hold all other things equal but make Y dry the match and X strike it, it still seems that D ought to kill Y, even though X does not now merely enable the killing.

Perhaps it might be argued that the contrast between causing and enabling is more important where those involved are less culpable. But this is doubtful. Consider the following:

20. See, e.g., Hart and Honoré, *Causation in the Law*; and Peter Menzies, "Causation in Context," in *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, ed. Huw Price and Richard Corry (Oxford: Oxford University Press, 2007), 191–223.

21. For the general difficulty of distinguishing metaphysical and linguistic questions in this context, see Eric Swanson, "Lessons from the Context Sensitivity of Causal Talk," *Journal of Philosophy* 107 (2010): 221–42.

Burned Car: V has been knocked unconscious by a kidnapper and is locked in the trunk of an old car. X and Y are about to set fire to the car for fun, not realizing that V is inside. However, they are both negligent about V's presence. If they set fire to the car, V will be killed. They agree that Y will pour petrol over the car. X will light the match. Z can prevent them from setting fire to the car only by shooting either X or Y in the leg.

Z seems to have no more reason to shoot X than Y, even though Y seems to enable the threat whereas X causes it. Overall, the contrast between enabling and directly causing does not seem morally significant, even if there is a robust way of distinguishing these things.

VII. LEWIS-STYLE INFLUENCE

Another view is that the causal contribution that one event makes to another depends on how sensitive the second event is to the first. David Lewis developed a counterfactual account of this idea. He argued that E1 causes O only if there is a substantial range of alterations on E1 that would alter O, and the degree of the causal contribution that E1 makes to O depends on that range.²² This proposal provides a plausible alternative explanation for why Y's conduct seems a more important cause of the audience being moved in *Performance* than X's conduct. A range of alterations in Y's performance gives rise to a range of alterations in the audience's reaction. This is not true to the same degree of X's recording. *Lewis-Style Influence* is the version of *Causal Contribution* that relies on this idea.

Before considering *Lewis-Style Influence*, I respond to an objection to Lewis's general view. As we will see, although this objection is forceful, it is not a decisive objection to *Lewis-Style Influence*. Lewis developed his account of causation to deal with the problem of preemption. Consider his case:

Rock: Suzy throws a rock, which hits a bottle, and the bottle smashes. Billy throws a rock, which arrives immediately afterward and flies through the space where the bottle was.

Had Suzy not thrown her rock, the bottle would nevertheless have been smashed, but the same thing is true of Billy's throw. The simplistic counterfactual view of causation described in Section III has no way to distinguish Suzy's throw from Billy's, but it is clear that Suzy's throw, not Billy's, causes the bottle to smash.

22. Lewis, "Causation as Influence," 191.

Lewis proposed that the causal difference between them is explained by the fact that small adjustments in Suzy's throw, but not Billy's, affect the smashing, and this makes Suzy's throw have much more causal power than Billy's. He writes,

Even if the throws are so much alike that removing Suzy's throw altogether would make little difference to the shattering, it is still true that altering Suzy's throw slightly while holding Billy's fixed would make a lot of difference to the shattering, but altering Billy's throw slightly while holding Suzy's fixed would not. Take an alteration in which Suzy's rock is heavier, or she throws a little sooner, or she aims at the neck of the bottle instead of the side. The shattering changes correspondingly. Make just the same alterations to Billy's preempted throw, and the shattering is (near enough) unchanged.²³

This response to the problem of preemption fails. To see why, note that these two things are consistent: (1) Suzy's rock actually hits the bottle, smashing it; and (2) in many nearby possible worlds where Suzy's throw is held constant but Billy's throw is altered, the time and manner of the smashing are altered. These things are true if small alterations in Billy's throw would result in his rock hitting the bottle rather than Suzy's. Where this is true, Lewis's theory implausibly implies that the causal potency of Billy's throw is increased.

First notice that if Suzy's throw occurs just before Billy's, small alterations in Billy's throw alter the time and manner of the smashing—those alterations that make Billy's throw occur before Suzy's. But the causal potency of Billy's throw is not greater in cases where it occurs just after Suzy's. If his rock does not hit, it is completely causally impotent. Furthermore, suppose that there is a powerful constant wind that would have been behind Billy's throw but for a van that drove behind Billy, very briefly interrupting the wind. In worlds where the wind is behind his throw, Billy's rock arrives first. In this variation, hold Suzy's throw fixed but make Billy's occur slightly later. The time and manner of the smashing alter, but again, this fact does not increase the causal potency of his actual throw.²⁴

One response is to supplement a counterfactual account of causation with an account of complete and incomplete causal chains. Billy's throw makes no causal contribution, on this view, because there is no complete causal chain between his throw and the smashing. This is consistent with the idea that the magnitude of the causal contribution that

23. *Ibid.*

24. For related objections to Lewis's view, see Jonathan Schaffer, "Causation, Influence, and Effluence," *Analysis* 61 (2001): 11–19.

one event makes to another depends in part on the pattern of counterfactual dependence that Lewis relies on.

To see this, suppose that E1 and E2 are each part of a complete causal chain resulting in O. Small alterations in E1 lead to small alterations in O, whereas small alterations in E2 lead to no alterations in O. Lewis's view, as amended, implies that E2 has less causal influence than E1. This nicely explains our causal judgments in *Performance* and *Torture*.

The related version of *Causal Contribution—Lewis-Style Influence*—is that these differences in causal contribution make a difference to liability. This may help to explain our intuitions in some cases. Consider Fabre's view that tightening the screws on a tank makes too small a causal contribution to the threat posed by the tank driver to make that person liable to be harmed. Small differences in the manner and time of the screw tightening make no difference to the manner and time of the deaths caused by combatants. In contrast, small differences in the manner and time of shooting make a large difference to the manner and time of death. This, it might be argued, explains why those who tighten screws are not liable to be killed whereas those who fire weapons are.

To evaluate this view, we must hold other factors that might influence liability constant. These include other facts that might influence the degree of a person's causal contribution, such as overdetermination, preemption, and remoteness, as well as independent factors that affect liability, such as culpability.

With this in mind, consider the following:²⁵

Button and Lever: X has a button in front of him, and Y has a lever in front of him. If both X presses his button and Y pulls his lever before a certain time, *t*, V will be killed. How and when X presses his button make no difference to how and when V will be killed if Y pulls his lever. How and when Y pulls his lever makes a significant difference to how and when V will be killed if X presses his button. X will press his button and Y will pull his lever, with the joint intention of killing V, if nothing is done. D can prevent this from occurring by killing either X or Y.

V's death counterfactually depends on both X's act and Y's act. However, X's act has less Lewis-style influence over the death than Y's. The manner and time of V's death are less sensitive to small alterations in X's act than small alterations in Y's. Some might think that this case casts further doubt on Lewis's view because it seems that X and Y will make an equal contribution to V's death if nothing is done.

25. This is a variation of a case offered by Schaffer in "Causation, Influence, and Effluence." In Schaffer's case, Y does nothing and X's act is sufficient to bring about the death.

But suppose that Y makes a greater causal contribution to the death than X if both act. *Lewis-Style Influence* implies that X is liable to suffer less harm than Y to avert the threat to V. Some might defend this view because it matters not only whether we live or die but also how and when we live or die. X is not in control of how V lives or dies, whereas Y is. But this difference cannot support *Lewis-Style Influence*. Suppose that of all of the times and modes of death that Y could bring about by pulling his lever, he brings about the one that is best for V. *Lewis-Style Influence* nevertheless implies that Y is liable to suffer greater harm than X because of worse alternatives that were available to him that he did not take. This is hard to believe.

Y may be liable to suffer greater harm if he causes V to die in a worse way than he could. But this is not best explained by *Lewis-Style Influence*, but rather by Y's increased culpability. If we eliminate this factor, for example, by making Y aware that pulling the lever will result in V's death, but making him nonculpably ignorant of the possibilities available to him, there seems to be no difference in their liabilities.

The same result is intuitive in cases of reduced culpability. Suppose that X and Y think that V is liable to be killed but negligently form this belief. Or suppose that X and Y are negligent about whether they will cause the death by pressing and pulling. It still does not seem very plausible that Y is liable to suffer greater harm than X. Overall, Lewis's view of the way in which causation is scalar is at best controversial, and even if it is right, it has no bearing on liability.

VIII. OVERDETERMINATION, PREEMPTION, AND CAUSAL CONTRIBUTION

A person can causally contribute to a threat without the threat, or any portion of it, counterfactually depending on her act. This is so where the whole threat is overdetermined. It has been argued that the degree of causal contribution that one event makes to another depends on whether the result is overdetermined, and to what degree.²⁶

McMahan seems to rely on something like this as an explanation for why voters are not typically liable to be killed if they vote for a government that starts an unjust war.²⁷ His explanation is not quite this. He suggests that it is the lack of capacity to influence their government that makes them nonliable. As voters together cause the election of the government that starts the war, these voters do play a causal role in the unjust war.²⁸ What McMahan probably means is that the actions of the gov-

26. See Braham and van Hees, "Degrees of Causation."

27. McMahan, *Killing in War*, 225; Fabre, *Cosmopolitan War*, 77.

28. See also McMahan, *Killing in War*, 217–18.

ernment do not counterfactually depend on the acts or omissions of any individual voter.

A. *Overdetermination*

Let us focus on the simplest overdetermination case first: E1 and E2 are each sufficient without the other, in conjunction with all events other than E1 and E2 that occur, to bring about an outcome, O. But if neither E1 nor E2 had occurred, O would not have occurred.

Here are two kinds of overdetermination case. In preemption cases, E1 causes O. E2 makes no causal contribution to O, but it would have caused O had E1 not occurred. For example, two tigers are chasing an antelope. The first tiger kills it. Had it not done so, the second tiger would have killed it. It would have survived had neither tiger been present. But only the first tiger causes its death, preempting the act of the second tiger.

In cases of simultaneous overdetermination, E1 and E2 each causally contribute to O. For example, two people push a boulder over a cliff, and each pushes sufficiently hard that they would have pushed the boulder off the cliff alone.²⁹ There is a question about whether E1 causes O in this case.³⁰ Let us leave this more or less semantic debate aside—E1 makes a contribution to the cause of O, so let us say that E1 makes a causal contribution to O.

Here is how overdetermination can come in degrees. Suppose that a set of events, {E}, causally contribute to O. Some subset of {E} would have been sufficient to cause O without all other members of {E}, given all other facts and events that occur. The larger the remainder, the more heavily overdetermined O is. For example, suppose that a group of n people together push a boulder off a cliff. Each exerts equal force, and $n - r$ people each pushing this hard would have been sufficient to get the boulder off the cliff. The larger r is, the more heavily overdetermined it is that the boulder is pushed off the cliff.

Now consider this simplified voting case:

Vote: There is a referendum on whether to prosecute an unjust war. A bare majority will result in war; anything short of a bare majority will not. There are 100 voters. X votes for war. No one abstains. X can have no other influence on whether the war occurs.

29. Carolina Sartorio pointed out to me that there is a further distinction between cases where each person in fact contributes part of the force necessary to get the boulder over the cliff but would have contributed the whole amount necessary had the other not acted, on the one hand, and cases where each person contributes the whole amount necessary, but more force than is necessary to get the boulder off the cliff is applied, on the other hand. I doubt that this distinction is morally important, so I leave it aside.

30. See Lewis, "Causation as Influence," 182.

If there are 51 votes for war, the war counterfactually depends on how X casts his vote.

Now compare two further variations. In *Vote 2*, 52 vote for war. In *Vote 3*, all 100 vote for war. The war counterfactually depends on X's vote in neither *Vote 2* nor *Vote 3*. But it might be argued that X's causal contribution to the war is greater in *Vote 2* than in *Vote 3* because the war is more heavily overdetermined in *Vote 3* than in *Vote 2*. One explanation is that more changes to the actual scenario would need to be made in order for X's vote to make a difference to whether the war occurs in *Vote 3* than *Vote 2*. Another is that the war counterfactually depends on X's vote in a wider range of nearby possible worlds in *Vote 2* than in *Vote 3*.³¹ *Overdetermination* is a version of *Causal Contribution* that claims that liability depends on the magnitude of overdetermination.

B. *Overdetermination, Permissibility, and Liability*

Here is one simple way in which overdetermination affects liability. Suppose that a person is certain to die whether or not one causally contributes to the death. This makes it easier to justify causally contributing to the death: given that one can do nothing to save the person's life, one is sometimes permitted to causally contribute to the death where doing so will achieve something valuable. And if one is justified in contributing to the death, one's liability is either reduced or negated.

It is not news that justifications affect liability. But inflicting harm can be justified because the harm is overdetermined, and this is one way in which a causal fact affects liability. To see this, consider the following:

Contribution to Save: X is certain to kill V whether or not Y contributes to the death. If Y contributes, he will save the life of another innocent person, Z. Y does this.

It would have been wrong for Y to act in a way that makes a difference to whether V lives or dies to save Z. But it seems permissible for Y to contribute to the death where doing so makes no difference to whether V lives or dies. Or, at least, this is true if Y acts only in order to save Z's life.³²

31. In more complex cases, many combinations of the events that occur would be sufficient to bring about the outcome. For example, suppose that E1, E2, E3, and E4 together cause E5. Now suppose that the following combinations, and only these, among the group, would have been sufficient to bring about E5: (E1, E2), (E1, E3), (E1, E4), (E2, E4). It is tempting to claim that E1 makes a greater causal contribution to E5 than any of the other events in virtue of the fact that E1 is a member of more sufficient sets. See Braham and van Hees, "Degrees of Causation." If what I say below about the simple case is right—that the degree of overdetermination is irrelevant to liability—we can set aside this more complex case.

32. I argue that this is one kind of case where intentions make a difference to permissibility in Victor Tadros, *The Ends of Harm: The Moral Foundations of Criminal Law* (Oxford:

If Y acts in order to save Z's life, Y is also liable to suffer much less harm to save V than X. So in this way, overdetermination makes a difference to liability. This, though, does not support *Overdetermination*. *Overdetermination* is the view that liability depends on the degree of a person's causal contribution to a threat, whereas the argument just presented is one that suggests that a person's liability sometimes depends on whether the threat counterfactually depends on the person's contribution. As I suggested earlier, simple counterfactual dependence does not come in degrees: either the person's conduct will make the relevant difference or not.

Contribution to Save would support *Overdetermination* if Y's liability depends on the degree to which V's death is overdetermined. Suppose that we introduce other people who would take X's place were X not present. I doubt that this would make a difference to Y's liability as such. Insofar as it seems to make a difference, this is best explained by the fact that we are more confident in concluding that V's death is certain to occur in cases where his death is more heavily overdetermined. If we make the relevant probabilities explicit, hold them constant, and vary the degree of overdetermination, I doubt that there is any difference in Y's liability.

C. *Culpable Overdetermination*

Still, the argument just offered might seem to support the view that voters are not liable to be harmed because they do not, and cannot, make a difference to whether anyone lives or dies. But we should not draw that conclusion too quickly. The argument just offered is that overdetermination affects the justification of making a causal contribution to a threat and that justification can affect liability. This does not show that overdetermination is relevant where making a contribution to a threat is unjustifiable, or where a justification could be offered, but the person is not motivated by the facts that make this true.

In order to make an assessment of these facts, we must hold other things equal, including culpability, the gravity of the threat, and the number of people who must be killed to avert the threat. The last consideration is especially important: where a threat is overdetermined by two acts, both acts must be prevented to avert the threat, and this may involve harming two people. Harming two people, even if they are culpable, is harder to justify than harming one.

Oxford University Press, 2011), chap. 7. For related discussion in the context of war, see Kai Draper, *War and Individual Rights* (Oxford: Oxford University Press, 2016), 91, 200–202. Adil Haque thinks that those who permissibly threaten others because the harm they threaten is overdetermined are liable to be killed to avert the threats they pose, though I am not sure whether he would hold this view in cases such as those outlined here. See Haque, *Law and Morality at War*, 72–73.

Here are two cases that do a decent job of holding these things equal:

Fire: X and Y are each about to start a fire that, if nothing is done, will kill V. The fire would not reach V without both fires being started. D can either kill both X and Y, averting the threat, or kill no one, resulting in V's death.

Fire 2: X and Y are each about to start a fire that, if nothing is done, will kill V. The fires will merge. Either X's fire or Y's fire would have been sufficient without the other fire, in conjunction with the remaining facts, to give rise to an identical threat to V. D can either kill both X and Y, averting the threat, or kill no one, resulting in V's death.

To simplify our discussion, focus only on X. X's act causally contributes to the threat to V in both cases. However, in *Fire*, the threat that the fire poses to V counterfactually depends on X's act, whereas in *Fire 2* it does not. Is there a difference in X's liability in these cases?

To assess the significance of a person's causal contribution by comparing cases, we need to hold their mental state constant. But if a person knows that the result is overdetermined, she also knows that she does not raise the probability of the outcome, and that may affect culpability. To neutralize this fact, suppose that X and Y have the same evidence and the same beliefs in both cases: they act in order to raise the probability of the outcome with the intention of killing V. When we consider this pair of cases in this way, there seems to be no difference in the liability of X and Y in *Fire* and *Fire 2*. Furthermore, X is clearly liable to be killed to avert the threat that he will otherwise pose to V in both cases.

If this is right, it is wrong to think that a person cannot be liable to be killed to avert a threat if that person did not, and cannot, make a difference to whether a threat occurs. In *Fire 2*, X cannot make such a difference. The same is true of Y. But they are both clearly liable to be killed to avert the threat they together cause.

Perhaps it might be argued that overdetermination does make a difference where there are many contributors, as there are in *Vote 3*. But this does not seem to make a difference, at least if we explicitly hold all else constant. If we introduce many other fire starters into *Fire 2* whom X does not know about, and if we are explicit that the probabilities of V dying are identical across different cases, I doubt that his liability to be harmed alters (though, of course, it may be more difficult to justify killing anyone, as killing more people would be required to avert the threat).

The same thing seems true in cases where the person's conduct is justifiable but was not justified. Conduct is justifiable but not justified

where the person had sufficient reason to act but did not act for that reason. For example, if I kill a person simply out of hate, but that person was in fact wrongly attacking me, my act was justifiable but not justified.

Suppose that in *Fire 2* Y starting the fire saves a further innocent person, but Y acts not in order to save this person, but only in order to kill V. Starting the fire is justifiable because Y starting the fire makes no difference to whether V lives or dies, but doing so saves a life. It is not justified because Y is not motivated by this fact. I doubt that those who act in a justifiable but unjustified way escape liability, or even have diminished liability, when compared with those whose acts are neither justifiable nor justified. Y's liability seems to depend on the reasons for which he acts, not on the reasons that were available to him to act.

To confirm this conclusion, consider this case I prepared in an earlier work:

*Poisoned Pipe.*³³ Boss offers a reward of £1000 to anyone who kills Victim. Two henchmen, X and Y, independently find different points in the water pipe leading to V's home. X puts sufficient poison in the pipe to kill V. At the same time, Y puts sufficient poison in the pipe to kill V. The poisons mix. Each sees what the other is doing. Neither can influence the behavior of the other. X's poison alone or Y's poison alone would have caused V to suffer a very slow and painful death. Together, their poison kills V swiftly and painlessly. When X and Y both act, V's death occurs at exactly the same time as it would have occurred had either X or Y acted alone. X and Y each acts only for the money.

In previous work, I argued that X and Y each act wrongly, even though X would have acted permissibly were she to have put poison in the pipe in order to make V's death less painful, and the same is true of Y. X and Y also seem liable to suffer a great deal of harm to save V's life. V seems clearly permitted to kill X and Y in order to save his life. This seems true even though X and Y each make V better off. The view that a person is not liable to be killed if they make no difference to whether a person lives or dies thus seems unsound; it even seems unsound where the act that makes no difference is justifiable.

The same thing seems true in cases of less culpability. There seems to be no difference in the liability of X and Y in the fire cases if they are merely negligent, for example. Overall, although overdetermination can affect liability by affecting whether a harm-causing act is justified, it does not seem to play a further role in grounding liability.

33. This is slightly modified from the version presented in Tadros, *Ends of Harm*, 159.

IX. PROBABILISTIC CAUSATION

One view that naturally implies that causal contribution is scalar is a probabilistic view of causation, for probability is clearly scalar.³⁴ Probabilistic causation is commonly defended as a theory of event-type causation (e.g., smoking causes cancer) rather than a theory of event-token causation (this act of smoking caused this person to get cancer).³⁵ Event-type theories do not plausibly have implications for liability—*Causal Contribution* is a view about the causal contribution that a person actually made to a threat, not about the causal power of the type of act that she performed.

Some have offered event-token probabilistic causal theories.³⁶ Such theories are hard to defend, given that some token causes decrease the probability of their effects and some events increase the probability of events they don't cause. For example, if I shoot at you and miss but another person shoots and hits, my shooting increased the probability of your death but made no causal contribution to your death. Conversely, suppose that I, a bad shot, shoot you dead. Had I not shot, a trained marksman would have shot at you, but my shooting ensures that he does not shoot regardless of whether I hit or miss. My shooting decreased the probability of your death without this reducing the causal power of my shot.³⁷

These difficulties may not be decisive. Perhaps probability is a component of a more complete causal theory. We might rely on the distinction between complete and incomplete causal chains introduced earlier to respond to the first problem. And a more complex view about the relationship between probability and preemption might resolve the second.

Probabilistic theories also differ with different kinds of probability or theories of probability. *Epistemic Probability* is a version of *Causal Contribution* that relies on epistemic probabilities. The metaphysical element of this view is in one way less controversial, because it is less controversial that epistemic probabilities exist. It is in another way more controversial. The epistemic probability of my act causing an event depends on the ev-

34. For a complex version of roughly this idea, see Alex Kaiserman, "Causal Contribution," *Proceedings of the Aristotelian Society* 116 (2016): 387–94.

35. For a defense of the view that different causal theories are needed for type and token causal claims, see Ellery Eells, *Probabilistic Causality* (Cambridge: Cambridge University Press, 1991).

36. See, for an overview and discussion, Christopher Hitchcock, "Probabilistic Causation," in *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta (Stanford, CA: Stanford University), <http://plato.stanford.edu/entries/causation-probabilistic/>.

37. Increasingly complex versions of probabilistic theories respond to concerns like this. See, e.g., Eells, *Probabilistic Causality*, chap. 6; Judea Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. (Cambridge: Cambridge University Press, 2009).

idence available to me. But, at least at first blush, causation does not seem to depend on evidence in this way. For example, the epistemic probability of my killing a person when I pull the trigger of a gun depends on my evidence that the gun is loaded, but this does not seem to affect whether my pulling the trigger causes a death.

If causation does depend on evidence, *Epistemic Probability* is roughly right. Here is why. The magnitude of a person's liability to be harmed depends on how culpable she is for posing a threat. The degree of a person's culpability for posing a threat sometimes depends on the magnitude of the epistemic risk that she takes in posing the threat. If the magnitude of a person's causal contribution to a threat depends on the magnitude of the epistemic risk that she takes in posing that threat, liability depends on facts that determine the magnitude of a person's causal contribution to the threat.

Even if *Epistemic Probability* is true, it adds little to our understanding of liability. The view that epistemic risk makes a difference to liability is both familiar and intuitive, regardless of beliefs about the relationship between epistemic risk and causation. We can safely bypass controversial metaphysical views about the relationship between epistemic risk and causation, as those who write about risk in the ethics of war currently do.³⁸

Objective Probability is a probabilistic version of *Causal Contribution* that claims that the probabilities relevant to both causation and liability are objective. This view is also metaphysically controversial, because it is controversial whether the probability of one token event causing another can be other than 0 or 1 in a deterministic world where all of the properties of the first token event are held constant.³⁹ For example, suppose that I flip a coin, which lands on heads. It is plausible that if the world is deterministic, the probability that this coin flip, with all of its properties, lands on heads is 1, and the probability that it lands on tails is 0.

Some might respond that the world might not be deterministic. But even if this is true, it is doubtful that liability depends on any kind of indeterminacy that makes determinism false, such as quantum indetermi-

38. See, e.g., Lazar, *Sparing Civilians*, chap. 4; Haque, *Law and Morality at War*, chaps. 5–8; Victor Tadros, "Uncertainties of War," in *Oxford Studies in Philosophy of Law*, ed. John Gardner, Leslie Green, and Brian Leiter, vol. 3 (Oxford: Oxford University Press, forthcoming).

39. David Lewis, for example, thought it obvious that it cannot be. See David Lewis, "Postscripts to 'A Subjectivist's Guide to Objective Chance,'" in *Philosophical Papers*, vol. 2 (Oxford: Oxford University Press, 1986), 118–21. For a more developed skeptical analysis, see Jonathan Schaffer, "Deterministic Chance?," *British Journal for the Philosophy of Science* 58 (2007): 113–40. Several attempts have been made to show that objective probability is compatible with determinism, but those working on probabilistic causation remain undecided. See, e.g., Branden Fitelson and Christopher Hitchcock, "Probabilistic Measures of Causal Strength," in *Causality in the Sciences*, ed. Phyllis McKay Illari, Federica Russon, and Jon Williamson (Oxford: Oxford University Press, 2011), 600–637.

nacy—is it hard to believe that when people lacked evidence for genuine indeterminacy, their judgments of liability ought to have been different from ours.

And even if there are objective probabilities, there is disagreement about whether such probabilities make a difference to the magnitude of causal contribution.⁴⁰ Suppose that they do. Could these probabilities be relevant to liability? To answer this question, we should consider pairs of cases where some causally contributing event increases the probability of its effect on any theory—that allows us to abstract from the complexities of different probabilistic theories of causation. We can do this by considering simple cases that do not involve confounding factors, such as overdetermination, and by simply stipulating objective probability.

Furthermore, notice this distraction: differences that a person makes to the objective probability of a threat occurring often track differences in the probability that killing the person will make to averting the threat. The second fact is often morally significant, but it provides no support for *Causal Contribution*.

We can eliminate this distraction, and when we do, objective probability does not seem to affect liability. One way to eliminate it is to consider cases where the fact that affected the likelihood of the outcome has passed. Compare the following cases:

Auto Kill 1: Sheila wishes to kill Jake, who is on the other side of a canyon, for no good reason. There is a safe bridge across the canyon, and Sheila knows that it is safe. She has a car that can be programmed to kill Jake. She programs it to cross the bridge and shoot him. She gets in the car because she wants to make sure that Jake is dead. It crosses the bridge, and it is about to shoot Jake. You can destroy the car with Sheila in it.

Auto Kill 2: Same as *Auto Kill 1*, except that there is a high objective probability that the bridge will collapse when the car crosses it. It looks to Sheila exactly the same as the bridge in *Auto Kill 1*. The car crosses the bridge, and it does not collapse.

There is a high objective probability of Sheila killing Jake by programming the car in *Auto Kill 1*, but only a low objective probability of her doing so in *Auto Kill 2*. But there is an equal chance of eliminating the threat by killing Sheila in both cases.

Suppose that differences in the objective probability of the threat being posed as a result of Sheila programming the car make a difference

40. See, for some discussion, Jon Williamson, “Probabilistic Theories,” in *The Oxford Handbook of Causation*, ed. Helen Beebe, Christopher Hitchcock, and Peter Menzies (Oxford: Oxford University Press, 2009), 185–212.

to Sheila's causal contribution in these cases. This seems to make little difference to Sheila's liability. She seems equally liable to be killed to save Jake's life in both cases. The same thing seems true if Sheila is less culpable, for example, if she is negligent about Jake's liability to be killed, or about whether her act will harm Jake.

Another way to eliminate distractions about the likelihood that harming the person will avert the threat is to compare cases where other things are equal but the probabilities are epistemic in one case and objective in the other. Compare the following:

Launch 1: There is a code to launch a missile that will pose a lethal unjust threat. Sandra knows the whole code except the last number. She is unsure whether that number is 0 or 1, and she has no more reason to believe one or the other. She has only one chance to enter the code. She is about to enter 0. Lucy has the same information as Sandra and can harm her to prevent her from doing this.

Launch 2: Same as *Launch 1*, except that, unknown to Sandra and Lucy, the last number is generated randomly by some indeterministic mechanism after the button is pressed, so that there is a .5 chance that her entering 0 will launch the missile.

In each case, the epistemic risk that Sandra will launch the missile, from both her perspective and Lucy's, is 50 percent. However, in the first case the objective probability that she will launch the missile is either 1 or 0. In the second case it is .5. The latter is true for all theories of objective probability, whether or not objective probability is compatible with determinism.

According to *Objective Probability*, if Sandra launches the missile, she makes a larger causal contribution to the threat in *Launch 1* than *Launch 2*. Obviously, if Sandra does not launch the missile, she makes no causal contribution to a threat. *Objective Probability* thus implies that Sandra is liable to suffer a greater degree of harm to avert the risk of posing a lethal threat in *Launch 1* than in *Launch 2*. Here is why. If Sandra gets the right number, she makes a larger causal contribution to the threat in *Launch 1* than in *Launch 2*. If Sandra gets the wrong number, she makes no causal contribution to the threat in either case. Hence, if Sandra is harmed, there is a .5 epistemic chance that she will be prevented from making a larger causal contribution to the threat in *Launch 1* than in *Launch 2* and a .5 epistemic chance that she will be prevented from making no causal contribution to the threat in either case. If causal contributions matter to liability, it naturally follows that there are stronger reasons to kill Sandra in *Launch 1* than in *Launch 2*. But it is not intuitive that there is such a difference. This provides further support for the view that *Objective Probability* is false.

To reinforce this conclusion, compare *Launch 2* with the following:

Launch 3: Same as *Launch 2*, except that the last number is randomly generated prior to Sandra pressing the button.

In *Launch 3*, Sandra pressing the right button is objectively certain to launch the missile. *Launch 3* is thus identical to *Launch 1* with respect to objective probability. *Objective Probability* thus implausibly implies that the permissibility of killing Sandra depends on whether the last number is randomly generated prior to or after her pressing the button.

Overall, even if the metaphysical claims that probabilistic views rest on are sound, either they add nothing to our existing understanding of liability, or they are false.

X. PROXIMITY

Proximity is a version of *Causal Contribution* that relies on the idea that the causal contribution that one event makes to another depends on how proximate these events are to each other. Voters, or those who tighten the screws on tanks, or those who manufacture weapons, are much earlier in the causal chain that leads to the threats posed in war than those who load and fire weapons. *Proximity* claims that these facts affect liability.

There might be both qualitative and quantitative aspects to proximity. The most familiar and important qualitative issue concerns wrongful intervention.⁴¹ Elsewhere, I argue that wrongful intervention does not itself make a moral difference, and I will not repeat that argument here.⁴² But even if it makes a difference in some cases, it does not plausibly make a difference where the first person intends the wrongful act of the second. In that case, the second person is not truly a wrongful intervener at all—their act is part of the first person's plan. That is often true in war, where noncombatants and combatants are part of a common enterprise.

Another possibility is that quantitative proximity to the threat makes a difference to one's liability to be harmed to avert that threat. Michael Moore, for example, claims that causation is not transitive—the fact that E1 causes E2 and E2 causes O does not imply that E1 causes O.⁴³ This is so, Moore believes, because whether E1 causes O depends on the number of events in the chain between E1 and O. It might also be claimed that

41. See, e.g., David Rodin, "The Myth of National Self-Defence," in *The Morality of Defensive War*, ed. Cécile Fabre and Seth Lazar (Oxford: Oxford University Press, 2014), 69–89, 82; Jeff McMahan, "Proportionate Defense," *Journal of Transnational Law and Policy* 23 (2013–14): 1–36, 36; Frowe, *Defensive Killing*, chap. 5.

42. See Victor Tadros, "Permissibility in a World of Wrongdoing," *Philosophy and Public Affairs* 44 (2016): 101–32.

43. See Moore, *Causation and Responsibility*, 155.

the causal contribution that one event makes to another depends on proximity.

This view is hard to vindicate. As Helen Beebee notes, events can be subdivided into further events. There may be no principled way to avoid infinite subdivisions.⁴⁴ For example, starting a car involves turning a key. Turning a key involves an infinite number of infinitesimally small turns. So starting a car involves an infinite number of events. We might rely on intuitive judgments about the length of causal chains—we do, after all, seem able confidently to make claims about proximity and remoteness. The problem of individuating events may well not be lethal to the idea of proximity and remoteness.

Even if this concern can be met, though, *Proximity* has counterintuitive implications. Where there is a chain of events that leads to a threat of harm and all other things are held equal, I doubt that it matters to a person's liability where in the chain the person's act occurs.⁴⁵

Consider the following:

Pass: a chain of combatants passes a weapon to the front line, where it will be used to kill an innocent person.

On the view under consideration, the magnitude of harm that those in the chain are liable to suffer depends on where in the chain they are, but this does not seem true. If there is a reason to harm those later in the chain rather than earlier, this is best explained by the fact that there is a greater likelihood that the munitions will get to the front line the farther down the chain they get.

Or consider this slightly different case:

Flood: One hundred people wish to cause a flood, which will devastate a village. To do so, they need to remove a flood-defense barrier. They can do so by firing rockets at it. If one hundred rockets hit the barrier, it will break, and the village will be flooded. Each of the one hundred people has a rocket. D can prevent the flood from occurring by killing any of the one hundred. He knows the order in which they will fire their rockets.

The firing of the first rocket is more remote from the village being flooded than the firing of the last rocket. The view that remoteness makes a difference to liability implies that D has a reason to kill the person who will fire his rocket last. But again, any such reason that D has is best explained by the fact that the probability of the threat occurring if D does not shoot a person increases as more rockets are fired. Even if D has a reason to

44. Beebee, "Legal Responsibility and Scalar Causation," 106–7.

45. See also McMahan, "Basis of Moral Liability," 396.

kill the last person in the chain, each person seems liable to be killed to avert the threat to the village where only killing this person can avert the threat.

And our judgments are no different if culpability is reduced. Make those who fire the rockets negligent about whether the people in the village are liable to be killed, or about whether the village will be flooded, and the view that D should harm the last person in the chain seems no more plausible.

XI. CONCLUSION

If their metaphysical claims can be vindicated, some versions of *Causal Contribution* seem true. For example, *Personhood* and *Epistemic Probability* both seem true. But they do not offer us something distinct from familiar ideas in moral philosophy. At most they clarify and deepen our understanding of these familiar ideas. Other versions of *Causal Contribution* seem false, both because the metaphysical ideas on which they rest are fishy and, more importantly, because they do not seem morally plausible. This is true of *Portion of the Cause*, *Enabling*, *Lewis-Style Influence*, and *Objective Probability*. Still other versions merit a more complex response. There is room for doubt about whether *Portion of the Threat* is true. And although *Overdetermination* is false, we have seen that overdetermination can affect justification and that it can, for that reason, affect liability.

Some unjust noncombatants, as well as some unjust combatants, are not liable to be killed to avert the threats they contribute to, and that may be in part because of the size of the causal contribution they make to these threats. But there is little support to be found here for a more general principle of noncombatant immunity.

In more recent work, Fabre suggests that a mix of causal ideas might do the job.⁴⁶ Given the familiar idea that contextual interaction is morally significant, we should not reject out of hand the idea that a causal fact has moral significance only when it interacts with some other causal fact. But I doubt that contextual interaction is significant in this case. At least, I am not sure how to make progress with the suggestion.

Some might respond that the arguments offered must be wrong because they conflict with the brute intuition that many noncombatants contribute too little to wars to make them liable to be killed.⁴⁷ I doubt that this is right. Our intuitions may be justified in part because causal facts often coincide with other facts that are morally significant, such as epistemic risk and the probability of success in averting threats. Furthermore, we have good reason not to rely too heavily on brute intuitions of this

46. See Fabre, *Cosmopolitan Peace*, 6.

47. Fabre continues to endorse *Causal Contribution* for roughly this reason. See *ibid.*

kind. Such intuitions are likely to be influenced by a very wide range of facts about noncombatants, such as facts to do with culpability. And they are likely to be influenced by the powerful, but recent, trend of seeing the targeting of noncombatants as abhorrent. However, this trend is best explained by abhorrence of killing noncombatants to influence governments by terrorizing the population. Abhorrence of this practice may well be justified, but it should not lead us to abhorrence of the targeting of noncombatants to prevent them from causing unjust threats.

It does not seem abhorrent to target many noncombatants who are wrongly involved in weapons production, or military training, or research and design, to prevent them from contributing to unjust wars where this will save an equal number of noncombatants who are completely innocent from the lethal threats they will otherwise face. Prior to careful reflection, the view that this is permissible is plausible enough, as is the view that it is wrong. Which view we should ultimately accept depends on our assessment of moral principles like *Causal Contribution*.