

The Worse Than Nothing Account of Harm
and the Preemption Problem¹

Abstract. Because harm is an important notion in ethics, it's worth investigating what it amounts to. The counterfactual comparative account of harm, commonly thought to be the most promising account of harm, analyzes harm by comparing what actually happened with what would have happened in some counterfactual situation. But it faces the preemption problem, a problem so serious that it has driven some to suggest we abandon the counterfactual comparative account and maybe even abandon the notion of harm altogether. This paper defends a version of the counterfactual comparative account that solves the preemption problem, a version called the "worse than nothing account." It says that you harm someone just in case you leave them worse off than if you'd done nothing at all.

Harm is an important notion in ethics and so it's worth investigating what it amounts to. One way of understanding harm is in terms of counterfactuals. Here's an example:

Something harms a person just in case it leaves them worse off than they would have been otherwise.

This sort of account is known as the "counterfactual comparative account" (CCA) of harm because it analyzes harm by comparing what actually happened with what would have happened in some counterfactual situation.²

¹Thanks for feedback to Josh DiPaolo, Nate Fulton, Gregory Glatz, Tino Petrocelli, Dan Rabinoff, Elena Rabinoff, Casey Woolwine, and some anonymous referees.

²For discussion of counterfactual accounts of harm, see e.g. [Boonin, 2014, 52-103], [Bradley, 2012], [Feinberg, 1984, 34], [Feit, 2015], [Feit, 2017], [Feldman, 1991], [Hanna, 2016], [Hanser, 2008], [Johansson and Risberg, 2019], [Klocksien, 2012], [Norcross, 2005], [Northcott, 2015], [Purves, 2016], [Purves, 2019], [Rabenberg, 2015], [Roberts, 2007], [Thomson, 2011].

The CCA is commonly thought to be the most promising sort of account of harm there is.³ Part of its appeal comes from its intuitiveness,⁴ part from its explanatory power,⁵ and part from its ability to dodge objections that other accounts face.⁶

That said, there are some problems that have been raised for the CCA. One of the most important ones concerns what has been called the “preemption problem”, a problem involving cases in which something bad happens that preempts something even worse from happening. Here is a case that illustrates the problem:

KNIGHT. Bobby Knight gets mad at a philosopher, Phil, in response to a perceived slight. So he chokes Phil. Luckily for Phil, Knight is in anger management. It’s taught him to better control his behavior. He choked Phil because he applied anger management techniques. If he hadn’t, he would have dismembered Phil. [Hanna, 2016, 263].⁷

The CCA seems to issue the wrong verdict on these sorts of cases. For instance, in the case at hand, it intuitively seems that Bobby Knight harmed Phil; after all, he choked him. But the CCA seems to yield the verdict that Knight did not harm Phil. After all, if Knight hadn’t choked Phil, he would have dismembered him, which is even worse.

The preemption problem is often viewed as one of the most serious problems facing the CCA, if not the most serious. For instance, Neil Feit writes:

³Neil Feit writes: “It is widely acknowledged that the most plausible account of the target notion of harm is the counterfactual comparative account” [Feit, 2015, 361]. See also [Bradley, 2012, 396], [Norcross, 2005, 150].

⁴For those who note its intuitiveness, see e.g. [Boonin, 2014, 52], [Feit, 2015, 362], [Feit, 2017, 810], [Johansson and Risberg, 2019, 351], [Klocksiam, 2012, 285].

⁵For those who emphasize its explanatory power, see e.g. [Boonin, 2014, 52], [Bradley, 2012, 397], [Feit, 2015, 362], [Hanser, 2008, 434-7], [Johansson and Risberg, 2019, 351], [Klocksiam, 2012, 285], [Purves, 2016, 69], [Purves, 2019, 2631].

⁶For criticism of alternative accounts of harm, see e.g. [Boonin, 2014, 56-102], [Bradley, 2012, 398-405], [Johansson and Risberg, 2019, 351-352], [Klocksiam, 2012, 287-8], [Rabenberg, 2015, 2-8], [Woollard, 2012].

⁷The case is originally due to Alastair Norcross [Norcross, 2005, 165-6]; I’m quoting a condensed version of it from Nathan Hanna. For discussion of it and similar cases, see e.g. [Boonin, 2014, 62], [Bradley, 2012, 407-8], [Feit, 2015, 363], [Hanser, 2008, 433-4], [Johansson and Risberg, 2019], [Klocksiam, 2012, 287], [Norcross, 2005, 165-6], [Northcott, 2015, 159], [Rabenberg, 2015, 10], [Roberts, 2007, 271], [Thomson, 2011, 446-7].

Let's use two examples to illustrate the problem of preemption. ... In each of these cases, there is an event that seems extremely harmful, and yet the counterfactual account implies that it is not harmful at all. Moreover, the counterfactual account seems to imply that no harmful event occurs in either of the cases. These are disastrous results [Feit, 2015, 363].

For another example of someone who takes this to be a serious problem, Ben Bradley writes, speaking of the preemption problem (and another problem, called the omission problem, which I'll be discussing later in this piece):

The problems about preemption and omission strike me as very serious [Bradley, 2012, 398].

Indeed, Bradley thinks they are sufficiently serious that he eventually concludes that the CCA should be abandoned and indeed that:

The best course of action is to avoid appealing to the notion of harm at all in our moral theorizing. [Bradley, 2012, 410-11].

Likewise, in a recent article, Jens Johansson and Olle Risberg write:

As more or less all participants in the debate agree, preemption cases pose a serious challenge to CCA. [Johansson and Risberg, 2019, 353]

The purpose of Johansson and Risberg's article is to critically examine recent attempts to address the preemption problem. Johansson and Risberg uncover serious flaws in all of these attempts, eventually concluding:

In this paper, we have discussed several ways of salvaging the counterfactual comparative theories of harm from the preemption problem. If our arguments have been sound, the preemption problem remains a serious difficulty for all these views. One may be tempted to add further modifications and extra clauses to the counterfactual account in order to try to avoid this objection, but we are skeptical of this enterprise. The preemption problem appears to reoccur in a systematic way for counterfactual theories of harm, even when they are modified with the precise aim of

avoiding it. This suggests (though of course it does not entail) that the proper response is abandonment of the counterfactual account, and not further modification [Johansson and Risberg, 2019, 364].

To summarize the last few pages: harm is an important notion in ethics and the CCA offers an attractive account of it. But the CCA faces a problem, the preemption problem, which is so serious that it has driven various people to suggest we abandon the CCA and maybe even abandon the notion of harm altogether.

Fortunately, I think the preemption problem for the CCA is solvable and in this paper I will be defending a version of the CCA that solves it, a version I call the “worse than nothing account.”

Before I give my account, let me pause to briefly clarify some terms. By “harm” I mean all-things-considered harm, as opposed to *prima facie* or *pro tanto* harm.⁸

Another clarification: some who wish to give accounts of harm focus on harmful *actions*,⁹ while others talk about harmful *events*, including events that do not involve actions.¹⁰ I will mostly focus on harmful actions, although I will discuss harmful events at some points in my second section.

To motivate my version of the CCA, let us return to the example in which Phil the philosopher is choked. We can imagine Bobby Knight trying to convince Phil that he wasn’t harmed. “After all”, Bobby says to Phil, “if you hadn’t been choked, you would have been dismembered, which is even worse.” It is natural for Phil to reply:

I agree that if you hadn’t choked me, you would have dismembered me, and that this would have been even worse. But that doesn’t mean I wasn’t harmed. Obviously I was harmed; you attacked me, which was scary and painful. I’d have been much better off if you’d just left me in peace.

⁸Most others also take the CCA to be an account of all-things-considered harm; see e.g. [Bradley, 2012, 393-4], [Feit, 2015, 361], [Feit, 2017, 810], [Hanna, 2016, 251], [Johansson and Risberg, 2019, 352], [Norcross, 2005, 150], [Purves, 2019, 2630]. By contrast, some focus on *pro tanto* harm, e.g. [Purves, 2016, 68].

⁹See e.g. [Boonin, 2014, 52], [Norcross, 2005, 150].

¹⁰See, e.g. [Bradley, 2012, 391], [Feit, 2015, 361], [Feit, 2017, 810], [Hanna, 2016, 251], [Johansson and Risberg, 2019, 351], [Klocksien, 2012, 285], [Northcott, 2015, 151], [Purves, 2016, 68], [Purves, 2019, 2631], [Rabenberg, 2015, 8-9].

The version of the CCA I will defend is inspired by this remark. In particular, it says that Knight’s choking Phil was harmful and it explains why by appealing to the fact that if Knight had just left Phil alone, Phil would have been better off. I will call this sort of account the “worse than nothing account” because it says that you harm someone just in case you leave them worse off than if you’d done nothing at all.¹¹

In the first section, I contrast the worse than nothing account with some other accounts that attempt to solve the preemption problem, showing how it avoids problems that these others face. In the second section, I look at some related issues to the preemption problem and what the worse than nothing account says about these. In particular, I look at what it can say about (i) the harmfulness of events (as opposed to the harmfulness of actions), (ii) cases in which someone fails to perform an action, and (iii) harms involving multiple agents.

1 Contrast with other comparative accounts of harm

The worse than nothing account agrees with other counterfactual comparative accounts that we should account for harm by comparing the actual situation with some other situation or situations, but it disagrees with the other accounts regarding which situation or situations is relevant.

¹¹I should note that the worse than nothing account is not entirely original; similar accounts have been discussed in the literature from time to time. (See e.g. [Bennett, 1995, 113-119], [Donagan, 1977], [Howard-Snyder, Summer 2002], [Kagan, 1989, 97-8].) But such accounts have not had many champions in recent years. As a result, if you look at discussions of the problem I am focused on – the preemption problem – you will notice the following: when people consider possible responses to these problems, they typically fail to mention the worse than nothing account. As a result it does not seem to have been noticed that the worse than nothing account can solve these problems. And, given the seriousness with which recent authors have taken these problems, the claim that the worse than nothing account can solve them seems like a noteworthy conclusion.

I should also note that I have stated the account at a general and intuitive level. I do think that it is worth developing further. In particular, I think it is worth developing precise accounts of each of the components of the worse than nothing account. But I also think it is valuable to show that, even at this fairly general and intuitive level, it already holds a lot of promise when it comes to solving some major problems facing CCA’s.

1.1 Contrast with the worse than the alternative account

For instance, most people who offer CCA's say that something harms someone if it leaves them worse off than if that thing hadn't occurred. On this view, the relevant counterfactual situation is the one in which the thing did not occur. For convenience, let me call this sort of account the worse than the alternative account.¹²

So, for example, regarding the situation in which Knight chokes Phil, those who endorse the worse than the alternative account would have us compare Knight's choking Phil with what would have happened had that thing not occurred. If Knight hadn't choked Phil, he would have pulled off Phil's arms. So those endorsing the worse than the alternative account would have us compare Knight's choking Phil with Knight's pulling off Phil's arms. Obviously, pulling off arms is worse. So the people endorsing the worse than the alternative account get the result that Knight did not harm Phil. But this is the wrong result.

The worse than nothing account, on the other hand, compares Knight's choking Phil with what would have happened if Knight had done nothing. If Knight had done nothing, Phil would have been free to go about his business in peace, and thus would have been better off. So the worse than nothing account gets the right result that Knight's choking Phil harmed Phil.

I have just been contrasting the worse than the alternative account with the worse than nothing account. In particular, I have emphasized that the former has a problem of the following sort: there are some cases – namely preemption ones – in which the worse than the alternative account fails to identify a harm, but the worse than nothing account correctly identifies the harm. This is a false negative. But there are also false positives, cases in which the worse than the alternative account says that a harm occurred even when it didn't.

For example, consider the following case:

KIND WORDS. I see that my friend is sad and say some kind words to her, cheering her up. If I hadn't said the kind words, I would have offered to take her out for a meal, which would have made her even happier.

¹²For formulations of this sort of account, see e.g. [Boonin, 2014, 52], [Bradley, 2012, 396], [Feit, 2015, 361], [Feit, 2017, 810], [Feldman, 1991, 215-6], [Hanna, 2016, 251], [Hanser, 2008, 423-4], [Johansson and Risberg, 2019, 351], [Klocksiam, 2012, 285], [Norcross, 2005, 150], [Purves, 2016, 69], [Rabenberg, 2015, 8].

The worse than the alternative account says that my saying kind words to my friend harmed her, because if I hadn't said these words, I would have offered to take her out for a meal, and she would have been even better off. But this is absurd; obviously, in saying kind words to my friend and cheering her up, I didn't harm her.¹³

The worse than nothing account gives the right result in this sort of case. It compares my saying kind words to my friend with my doing nothing. If I had done nothing, my friend would have been worse off. The worse than nothing account thus delivers the correct result: my saying kind words to my friend did not harm my friend.

1.2 Contrast with the “fundamental intentions” account

Nathan Hanna has recently offered a counterfactual comparative account on which we contrast what actually happened with what would have happened if the subject hadn't performed an action of the same type, writing:

An agent's free act is overall harmful for a subject S iff S is better off in the closest possible world where the agent doesn't freely perform an act of that type than S is as a result of the act. [Hanna, 2016, 266].

What does it mean to perform an act of the same “type”. By this, Hanna means one with the same “fundamental intention”. He writes:

For any free act that an agent performs, the fundamental intentions that she acts on constitute types under which we can classify the act for the purpose of harm assessment. [Hanna, 2016, 267].

For example, in the Knight case, Hanna says that Knight's fundamental intention is to hurt Phil. This means that on Hanna's account, the relevant contrast case is one in which Knight had acted without intending to hurt Phil. If he had acted without this intention, he presumably would have just left Phil alone. So Hanna gets the result that Knight harmed Phil, seeing as if he had acted with a different fundamental intention, Phil would have been better off.

This account is a little difficult to evaluate because the notion of a “fundamental intention” is somewhat obscure. Even so, it seems that there are

¹³For a somewhat related point, see [Northcott, 2015, 153].

counterexamples to Hanna's account. For instance, suppose we modify the Knight case to say that Knight has the delusional belief that clobbering someone with a baseball bat will help them be a better person. And let us suppose that if Knight hadn't tried to hurt Phil, he would have tried to help him by clobbering him with a baseball bat, and that this would have been even worse for Phil. In such a case, Hanna's account delivers the result that Knight didn't harm Phil by choking him, because if he had acted with a different fundamental intention, Phil would have ended up even worse off.

But this is the wrong result; Knight's choking Phil harmed Phil, and this is so even if it's the case that if Knight hadn't intended to hurt Phil, he would have clobbered him with a baseball bat.¹⁴ By contrast, the worse than nothing account has no problem handling this case. It says, correctly, that even in this modified case, Knight harms Phil; his choking him is worse than just doing nothing to him.

1.3 Contrast with the state of affairs account

Another sort of counterfactual comparative account is Fred Feldman's account which focuses on states of affairs. Feldman introduces it in the context of discussing a case of Jeff McMahan's involving the charge of the Light Brigade. I won't get into the details of this case, but rather just apply what he says to the Knight case.

Feldman's idea is to distinguish several states of affairs in this sort of case, and to say that some of them are harmful and others are not [Feldman, 1991, 225-6]. So, for example, regarding the Knight case, we could distinguish Knight's choking Phil from Knight's attacking Phil. We could then say that the choking is not harmful – if Knight hadn't choked Phil, he would have pulled off his arms, which is even worse. But we could say that Knight's attacking Phil was harmful; if Knight hadn't attacked Phil, he would have left him in peace, which would have been better.

One complaint that has been offered against this sort of response to the preemption problem is that Feldman has changed the subject. Many wish for an account of when concrete particulars – be they actions or events – are harmful, but instead of offering such an account, he's discussing the circumstances in which certain abstract objects – states of affairs – are harmful.¹⁵

¹⁴For a similar criticism, see [Johansson and Risberg, 2019, 360-364].

¹⁵For those who say this, see e.g. [Feit, 2015, 364], [McMahan, 2002, 120].

That said, this problem seems solvable. As Neil Feit notes, Feldman, if he wishes to do so, could modify his account so as to turn it into an account of the harmfulness of events or actions. In particular, he could switch from talking about states of affairs to talking about their instantiations at particular times and places. So, for example, he could talk about that concrete particular which is the (actual) spatiotemporal instantiation of Knight's choking Phil at the time in question and likewise he could talk about that concrete particular that is the (actual) spatiotemporal instantiation of Knight's attacking Phil at the time in question. He could declare these both to be actions and say that the first was not a harmful action and the second was [Feit, 2015, 364].¹⁶

Feldman's view, once modified so as to be a view regarding the harmfulness of actions, has some attractions with regards to the preemption problem. In particular, his view is able to deliver the result that Phil is harmed when Bobby Knight attacks him. After all, he is attacked and this is a harmful action. This is far better for Feldman than if he had to say that Phil was not harmed at all.

That said, Feldman's account faces an important problem. In particular, Feldman cannot say that the action of Knight's choking Phil is a harmful action. It seems better to be able to say – as the worse than nothing account does – that this is a harmful action.

To develop the idea further: not only do we want to say that Phil is harmed, we also want to say that the action of Knight's choking Phil is an action that harms Phil. To bring this out, let us look at the original version of the Knight case offered by Norcross. Norcross writes:

Bobby Knight, turning as red as his sweater, seizes the hapless philosopher around the neck and chokes her violently, while

¹⁶To do so, he would have to be a “multiplier” and think of these as different events, otherwise he'd get the bizarre result that one event was both harmful and not harmful. This fact that Feldman, in order to incorporate this strategy, would have to endorse a particular controversial view of the metaphysics of events, namely that of multipliers, is an advantage for the worse than nothing account over Feldman's account. Someone endorsing the worse than nothing account can remain uncommitted about the metaphysics of events, at least insofar as the unifier-multiplier debate is concerned while Feldman, if he were to take this strategy, would not be able to. (As I mentioned above, Neil Feit was the first to note that Feldman can turn his account of the harmfulness of states of affairs into an account of the harmfulness of events. But Neil Feit himself seems to be a unifier and, for this reason, criticizes the modified account, which only a multiplier can endorse.) [Feit, 2015, 364].

screaming obscenities. By the time Bobby Knight has been dragged away, the philosopher has suffered a partially crushed windpipe and sustained permanent damage to her voicebox. [Norcross, 2005, 165-6].

It seems clear, in considering this case, that the action of Knight's choking the philosopher is a serious harm. It seems deeply counterintuitive to have to claim that this action is not a harmful one. And thus the worse than nothing account has a clear advantage over Feldman's account.

It also has an advantage in the false positives case. In particular, recall the case in which my friend is sad and so I say kind words to my friend, thus cheering her up. Recall that if I hadn't said the words, I would have taken her out to lunch, which would have made her even happier. Feldman's account has to say that my saying kind words harmed my friend. But the worse than nothing account does not have to say this.

1.4 Contrast with the contextualist account

Another sort of CCA is a contextualist one. Such an account says that what counts as harm depends on conversational context.¹⁷

One way to motivate such an account is to appeal to the following claim: part of the semantic content of utterances involving counterfactuals is contributed by context. As Justin Klocksien, who defends a contextualist version of the CCA, writes: "... counterfactuals are highly sensitive to context, and there may be no single way things would have been if things had gone otherwise" [Klocksien, 2012, 289].

This sort of account might be thought to be able to handle the preemption problem as follows: it can say that the whether utterances regarding what Knight would have done if he hadn't choked Phil express true or false propositions depends, in part, upon the context in which they were uttered. Maybe in one context, the following utterance expresses a true proposition: "what he would have done is to leave Phil alone", while in another context, the following utterance expresses a true proposition: "what he would have done is to pull off his arms."

I should note that a contextualist account like Klocksien's has certain attractions. In particular, they can affirm – at least relative to certain contexts

¹⁷For those who endorse it, see e.g. [Klocksien, 2012, 289], [Norcross, 2005, 167-72], and [Northcott, 2015].

– that “Knight’s choking Phil harmed Phil” expresses a true proposition. In particular, according to the contextualist, there are some contexts in which the utterance “if Knight hadn’t choked Phil, he would have left him alone” expresses a true proposition. And in these contexts, the utterance “Knight harmed Phil” also expresses a true proposition.

That said, contextualists like Klockslem cannot handle cases in which it is made explicit, in the relevant context, that if Knight hadn’t choked Phil, he would have pulled off his arms. So, for example, when I presented case initially, I made it explicit that if Knight hadn’t choked Phil, he would have pulled off his arms. Nonetheless, in the context in which I initially presented the case, it still seemed as though Knight’s choking Phil harmed him. A contextualist like Klockslem cannot accommodate this judgment.¹⁸ But of course, the worse than nothing account can accommodate this judgment by noting that Knight’s choking Phil was worse than nothing.

This criticism applies to Klockslem because he is explicit regarding the factors that serve to make an alternative salient. There are other contextualists who are somewhat less explicit about this. For instance, Alastair Norcross offers a contextualist account in terms of “appropriate alternatives”. He does not fully explain what counts as an appropriate alternative, writing “I don’t here have the time (or the inclination) to give a detailed account of how conversational context determines the appropriate alternative.” [Norcross, 2005, 171].

This account nonetheless faces difficulties regarding the KNIGHT case. In particular, it faces a dilemma. Either, it has to say that, in the context in which I initially presented the case, the alternative in which Knight rips off Phil’s arm is an appropriate alternative. If so, it faces the same issues that Klockslem’s account does, viz. that it gives the wrong judgment about the case, namely that Knight did not harm Phil. The other option is to say that Knight’s ripping off the arm is not an appropriate alternative. If so, this account faces an explanatory burden: it has to explain why this alternative does not count as “appropriate”.

In addition to these problems, I should briefly note another problem with these contextualist accounts, a problem that those who offer such accounts acknowledge.¹⁹ As I noted earlier, many are inclined to give harm an impor-

¹⁸Others who note this criticism of contextualist CCA’s include: [Feit, 2017], [Johansson and Risberg, 2019, 355], [Purves, 2019].

¹⁹See e.g. [Norcross, 2005, 172].

tant place in ethics. But if what counts as harmful depends on conversational context, then it is hard to see how harm could have this important place. After all, it is doubtful that what we morally ought to do depends on conversational context.

1.5 Contrast with the contrastivist account

Let me discuss one more sort of account, namely the contrastivist account.²⁰ This sort of account says that harm is a relation between two actions. That is, actions aren't harmful simpliciter, rather they are harmful as compared with other actions. So, for example, Knight's choking Phil is harmful as compared with his leaving Phil alone, but not harmful as compared with his pulling off Phil's arms.

This sort of account might be thought to be able to handle the preemption problem as follows: it can say that Knight's choking Phil is harmful – at least relative to his leaving Phil alone. It thus avoids the problems facing the two accounts I mostly recently discussed. Feldman's account couldn't say that the choking was harmful, and the contextualist account had to say that in some contexts, it wasn't harmful. But the contrastivist account can say that in every context it is harmful, at least when compared with Knight's leaving Phil alone.

The problem for this sort of account is the following: for just about any action, there's a contrast on which it's a harm and a contrast on which it's not.²¹ So, for example, consider my donating a hundred dollars to a needy individual. Normally, this would not be viewed as harming the individual. But the contrastivist account cannot avoid saying that it is a harm, at least relative to certain possibilities. After all, I could have given two hundred dollars and the action of my only giving one hundred was a harm as compared with this.

Likewise, suppose I murder someone in a gruesome fashion. This seems like a harm. But it's not as bad as murdering them in an even more gruesome fashion. So the contrastivist also has to say I benefitted them, at least relative to this possibility.

Again, the worse than nothing account does not have to say these things. When I give someone a hundred dollars, as opposed to two, I do not harm them; a hundred dollars is better than nothing. Likewise, in murdering them

²⁰For discussion, see e.g. [Bradley, 2012, 407-8].

²¹For discussion, see e.g. [Bradley, 2012, 408].

in a gruesome fashion, I have harmed them because I've treated them worse than if I had just left them in peace.

Let me briefly summarize this section. As I noted, the worse than nothing account agrees with other CCA's that we should compare the actual situation with some counterfactual situation, but it disagrees regarding the question: what is the relevant counterfactual situation (or situations) we should compare the actual situation to? In doing a sweep of other sorts of CCA's defended in the literature, we have seen a large number of proposals. And for every proposal, we have seen some problems with it, problems that the worse than nothing account avoids.

2 Harmful events, multiple agents, and the omission problem

I have discussed the advantages of the worse than nothing account when it comes to the preemption problem. In this section, I want to briefly discuss three further issues. One concerns the fact that my account focuses on harmful actions, rather than harmful events. The second concerns what my account has to say regarding another important problem for counterfactual comparative accounts – the omission problem. The third concerns what my account says about harms involving multiple agents.

These topics are all rather complicated and I do not have the space here to fully discuss each and every aspect of them. Nonetheless, I hope that briefly discussing these topics will help forestall the worry that there is a decisive objection to the worse than nothing account lurking in one of these areas. More positively, I hope that the discussion illustrates some fruitful implications of the worse than nothing account.

2.1 Harmful events vs. harmful actions

As I indicated earlier in this paper, I am focused on offering an account of harmful actions, as opposed to harmful events.²² But some have chosen to offer accounts of harmful events instead. What would I say if they were to challenge my decision? That is, what would I say if they asked me the following question: why is it acceptable to focus on offering an account of harmful actions, as opposed to harmful events more generally?

²²I also noted that I am not the only one who does so. See e.g. [Boonin, 2014, 52], [Norcross, 2005, 150].

The first point I would make is that an account of harmful actions will provide much, if not all, of what we need to do ethics.²³ If we look at the contexts in which harm is discussed in ethics – contexts such as Mill’s harm principle, Ross’s views about prima facie duties and harm, and so on – we see that what is almost always being discussed is harmful action.

Even if this is right, a lingering worry remains. Perhaps, in giving an account of harmful actions, as opposed to harmful events, I will be failing to offer a unified account of harm. For instance, as Ben Bradley writes:

Often discussions of harm focus on actions performed by people. This is understandable, since it is largely in virtue of harm’s role in explaining the moral wrongness of actions that we are interested in it, and it seems that only people perform acts that are morally wrong. But acts performed by non-people, like cougars, can be harmful too. And many other sorts of events besides actions are harmful too, like explosions and earthquakes. They seem harmful in the same way that actions are harmful. An acceptable analysis of harm should allow for this. [Bradley, 2012, 394-5].

Let me start my response by saying that while it seems true to Bradley that non-actions are harmful in the same way that actions are, I am not completely convinced. First, a conciliatory point: I should acknowledge that there are some ways in which they are similar. For example, both earthquakes and the actions of persons can cause a decrease in someone’s well-being. Nonetheless, I think it is worth emphasizing that there are also important differences. For example, I can morally evaluate a person’s action but not that of an earthquake.

Because of this sort of difference, it seems to me to be an open possibility that when we speak of earthquake “harming” people, our talk is importantly different from when we talk of humans “harming” people. Compare: it makes sense to say that I blame the earthquake for my mother’s death. But this sort of blame seems to me to be importantly different from my blaming a killer for my mother’s death.

Thanks to these differences, one possibility is that when we talk about a earthquake being harmful, or when we talk about blaming a earthquake for my mother’s death, we are anthropomorphizing the earthquake, and thereby speaking somewhat metaphorically. Another related possibility is that there

²³A similar point is made at [Bradley, 2012, 394].

are certain senses of “harm” and “blame” that apply to harmful action but do not apply otherwise and that there are other senses of “harm” or “blame” that do apply to events that are not actions.

All that said, it seems to me an open possibility that Bradley is right and that events without actions can be harmful in the same way that actions are harmful. If this is indeed the case, I do not think this is a problem for the worse than nothing account. There are ways to expand the account to accommodate events without actions.

One sort of approach would be to expand the worse than nothing account beyond actions to what I will call “doings”. In short, the idea is that someone or something can “do” something, even if the thing it does is not an action. For example, if an earthquake disturbed my poodle, then this is something it “did”, even though its disturbing him was not an action.

Using this idea, we can expand the worse than nothing account to say that something harms someone if it leaves them worse off than if it had done nothing to them at all. For example, we can say an earthquake harms someone if it leaves them worse off than if it had done nothing to them at all. (What does it mean for an earthquake to do nothing to someone? One easy way for it to do nothing to someone is for it not to affect them in any way.)

Let me say a little more about what it means for a non-agent (like an earthquake) to do something.²⁴

First, the sense of “doing” I have in mind here does not require intentions. Indeed, even in the actions case, I do not think intentions are required in order to harm. So, for example, I can accidentally step on someone’s foot and thereby harm them even if my action was unintentional.

Second, as with human actions, there is a natural distinction between doing and allowing. So, for example, suppose there is a boulder that breaks free and rolls down a mountain, hitting a person and causing them injury. Suppose a park ranger could have stopped the boulder, but doesn’t. In this case, the park ranger don’t do harm to the person, but rather allows harm to be done to them. Likewise, we can imagine that there is a dead tree alongside the path of the boulder. If the tree had fallen over, it would have stopped the boulder, but it doesn’t fall over. In this case, just as with the park ranger, the tree, by not falling over, allows the boulder to harm the person.

²⁴Thanks to an anonymous referee for pressing me to say more here.

Third, just as with actions, we can distinguish enablers from doers.²⁵ So, for example, suppose the park ranger moves a dead tree off a path. Later on, a boulder rolls down the path and hits a person, causing them injury. While the park ranger didn't do the harming, she enabled it to occur. In a similar vein, a natural occurrence can enable harm without doing the harming. So, for instance, suppose that an earthquake shifts the ground slightly, causing the dead tree to move off the path. Sometime later, the boulder rolls down the cleared path, hitting the person. In this case, the earthquake enabled the harming, while the boulder did the harming.

As these examples bring out, a number of distinctions apply in the case of non-agential harms. For instance, as I have suggested, it is natural to distinguish between (i) cases in which a non-agent does something (such as the boulder rolling into the person) (ii) causes in which a non-agent allows something (such as the dead tree not falling, and thus allowing the boulder to roll past) and (iii) cases in which a non-agent enables something (such as the earthquake shifting the dead tree, and thus enables the boulder to roll past). These distinctions are important to respect, not merely for the worse than nothing account, but rather for any account that wants to match intuitions.

How shall we settle which, if any of these categories, a particular case falls into? As I've implicitly suggested, I find it helpful to ask: what if a human had done the same thing? This provides a useful heuristic that, along with our intuitions, can help settle cases of when non-agents count as doing things and what sorts of doings they are.

In sum, then it does not seem to be a problem that I focused on giving an account of the harmfulness of actions. First, actions are what matter most, ethically speaking. Furthermore, it is not clear that non-actions harm in the same way that actions do. Finally, the worse than nothing account can be expanded to handle cases of harmful non-actions.

2.2 The omission problem

Next, let me turn to the discussion of a problem closely related to the preemption problem, which is typically called "the omission problem. The omission problem involves correctly distinguishing harms from failures to act. Here, for example, is a case from Ben Bradley:

CLUBS. Suppose Batman purchases a set of golf clubs with the

²⁵For more discussion of enablers, see e.g. [Dancy, 2004, 38-43].

intention of giving them to Robin, which would have made Robin happy. Batman tells the Joker about his intentions. The Joker says to Batman, “why not keep them for yourself?” Batman is persuaded. He keeps the golf clubs. [Bradley, 2012, 397].²⁶

The CCA seems to issue the wrong verdict on these sorts of cases. In particular, it seems to deliver the result that Batman harmed Robin. After all, if Batman hadn’t kept the clubs, he would have given them to Robin, in which case Robin would have been better off. But this isn’t the result that matches intuition – or at least the intuitions of certain people. To these people, it intuitively seems that Batman didn’t harm Robin. Rather, he merely failed to benefit him.

How does the worse than nothing account distinguish between harms and failures to benefit? Recall that the worse than nothing account says that you harm someone just in case they end up the same as if you’d done nothing at all. This offers a natural contrast with failures to benefit: you fail to benefit them if they end up the same as if you’d done nothing at all.

This account delivers the correct result the case I just discussed. In particular, the worse than nothing account delivers the correct result that Batman did not harm Robin by failing to give him the clubs. This is because failing to do so is no worse than doing nothing at all.

Some other defenders of CCA, in responding to this sort of case, have suggested a different sort of response. In particular, they say that in this case, in not giving the clubs to Robin, Batman fails to act. And this is why his not giving the clubs to Robin does not count as a harm.²⁷

But this sort of response doesn’t handle simple variants of the case. So for instance, suppose that, as Batman is trying to decide what to do with the clubs, the Joker convinces Batman to give the clubs to the Joker, as opposed to Robin. In this case, Batman does a thing – he gives the clubs to the Joker. But still, it doesn’t seem as if he harms Robin. After all, he leaves him no worse off than if he’d done nothing to him at all.²⁸ This variant can be

²⁶For discussion of this and other related examples, see e.g. [Boonin, 2014, 53 n.2], [Feit, 2015, 367], [Feit, 2017], [Hanna, 2016, 252], [Hanser, 2008, 427-9], [Johansson and Risberg, 2019, 352], [Klocksiam, 2012, 286-7], [Purves, 2016, 81].

²⁷See e.g. [Boonin, 2014, 53 n.2], [Feit, 2015, 384-5], [Feit, 2017, 4], [Hanna, 2016, 252], [Hanser, 2008, 427], [Johansson and Risberg, 2019, 352]. Thanks to an anonymous referee for pressing me to discuss this response.

²⁸For a similar criticism, see [Feit, 2017, 812-813].

easily captured by the worse than nothing account. As I just noted, in this case, in giving the Joker the clubs, Batman leaves Robin no worse off than if he'd done nothing at all. Because of this, the worse than nothing account correctly delivers the verdict that Batman did not harm Robin. Meanwhile, other accounts falter. For instance, the worse than the alternative account says that Batman harms Robin in giving the Joker the clubs because if he hadn't given the clubs to the Joker, he would have given them to Robin, and Robin would have ended up better off.

What if the Joker's intervention comes after Batman has already acted on the intention of giving the clubs to Robin? So for instance, what if Batman gives the clubs to Robin and then the Joker convinces him to take them back, and so he does so? In this case, intuitions are more conflicted. On the one hand, it's hard to say that Batman harmed Robin; Robin started the story with no clubs and he ended it with no clubs, so he isn't worse off than at the start. On the other hand, it seems that in taking back the clubs Batman took away a thing that was – however briefly – Robin's, and thereby harmed him.

Here, I think that the worse than nothing account can accommodate our conflicting intuitions. In particular, it can say that if we focus on Batman's actions as a whole – his giving and then taking away the clubs – we see that he left Robin no worse off than if he'd done nothing, and thus he didn't harm him. But if we focus on Batman's final action – his taking away the clubs – we see that he left Robin worse off than he was previously, and thus harmed him.²⁹

I have just discussed some cases in which the worse than nothing account seems to correctly distinguish between harm and failure to benefit and thus to solve the omission problem. However, there are some other cases that seem to prove more problematic.

One sort of case that may seem to prove more problematic involves neglect. For example, suppose that Bob adopts a dog and then neglects it. The worse than nothing account seems to classify this neglect as a failure to benefit. After all, seemingly if Bob had done nothing to the dog, this would have been no worse than neglecting it. Indeed, neglecting the dog is a way of doing nothing to it. But this claim that Bob does not harm the dog seems

²⁹Thanks to an anonymous referee for suggesting this response. As with the group action cases discussed earlier in the paper, subtle questions remain about how, and when, an action being part of a group of actions that is harmful (or beneficial) affects whether it itself is harmful (or beneficial).

to fly in the face of intuition. Neglect of a dog seems like a harm, or, at the very least, like a very bad thing to do.

My response is as follows: in attempting to distinguish harm from failure to benefit, I have implicitly been assuming that this difference is morally significant. Otherwise, it wouldn't matter whether an account could distinguish harm from failure to benefit. But in assuming that this difference is morally significant, I do not have to hold that failures to benefit are never bad things. Nor do I have to hold that every harm is more serious than every failure to benefit. Rather, it is enough for me to hold that, other things being equal, it is worse to harm than it is to fail to benefit.³⁰

It is good that I do not have to hold these stronger claims. This is because they are very dubious. For example, consider a case in which Caroline accidentally steps on a stranger's foot, and causes him a minor amount of pain. Contrast this with a case in which David is walking by a pond and sees a child drowning and does nothing because he does not care. It seems as though David's action is very bad and far worse than Caroline's action, even though his action is a failure to benefit and her's is a harm.

In sum, then, I do not have to hold that every harm is worse than every failure to benefit. This allows me the following option: I can say that cases of neglect can be really bad, even if they involve a failure to benefit, rather than a harm. This claim is perfectly consistent with the claim that, other things being equal, failing to benefit is worse than harming. After all, there are many available explanations of why neglect is so bad – why it is one of the worst sorts of failure to benefit. For instance, it typically involves a lot of suffering, it typically involves a case in which someone has taken on an obligation to benefit and then failed to deliver, and so on.

2.3 Actions involving multiple agents

I have so far focused on cases involving a single agent. In such cases, someone does something and if they hadn't, they would have done something else. A natural question arises: what about cases involving multiple agents?

One such sort of case has the following structure: someone does something and if they hadn't, someone else would have done something which would have been even worse. Here, for example, is such a case: I see a pedestrian who is about to be killed by a drunk driver, and I pull her out of the way,

³⁰Here what I say is similar to [Purves, 2019, 2651].

thus dislocating her shoulder. While the pain hurts her, it leaves her better off than she would have been if she had been hit and killed by the driver.

This case is importantly different from the case of Bobby Knight. In particular, if I hadn't pulled her out of the way, it's not as if I would have done something else to her. Rather, it's someone else – the drunk driver – that would have done something else to her.

There are some other important differences as well. In particular, in the Bobby Knight case, Bobby Knight harms the philosopher. But in this case, it seems that I did not harm the pedestrian – at least not in the *ultima facie* sense of harm. Rather, I benefitted her; I saved her life!

What does the worse than nothing account say about this case? Recall that the worse than nothing account says, in general, that I harm someone if I leave them worse off than if I had done nothing at all. In this case, I do not leave the pedestrian worse off than if I had done nothing at all. On the contrary, I leave the pedestrian better off than if I had done nothing at all; if I'd done nothing at all, the pedestrian would be dead. Thus, the worse than nothing account says that I did not harm her, which is the intuitively correct result.

Dislocating a shoulder is pretty bad, but there are things that one can do that are even worse. So, for example, suppose that to save the pedestrian, I have to push her so roughly that she falls, breaking both her legs. Even though this is bad, it's not as bad as dying. And so the worse than nothing account says I still haven't harmed her. And intuition agrees. Her breaking her legs is less bad than her dying.

What about the extreme and tragic case – the case in which, in trying to save her life by pushing her out of the way, I end up causing her to die in some other way? Suppose, for example, that I push her out of the way of the oncoming car, without realizing that this will cause her to be hit and killed by a car coming in the opposite direction? In such a case, it seems that I have neither benefitted her nor harmed her, and the worse than nothing account agrees.

Let me quickly note a contrasting view and what is wrong with it. Someone who hadn't been considering the sort of case I just presented might initially be drawn to the following view: if someone is about to be killed and I kill them in exactly the same way, then I harm them a great deal. But this seems like a bad view, as we can see when asking what we should say about a case in which someone is about to be killed and I kill them in a way that is slightly less painful. For it seems that if someone is about to be killed and

I kill them in a way that is slightly less painful, then I do not harm them at all. (Hence the moral acceptability, according to many philosophers, of assisted suicide.) And it seems very strange to say that if I kill someone who was about to be killed in exactly the same sort of way, I cause them a great deal of harm, whereas if I kill them in a way that is slightly less painful, I cause them no harm at all.

One fact that these cases bring out is that if people are in extremely bad circumstances, one can benefit them even while doing things to them that are pretty unpleasant. So, for example, suppose that on a lonely road in the middle of winter I come across someone with car problems who is freezing to death. Suppose that I offer to drive them to safety if they promise me an enormous amount of money and suppose further that, quite desperate, they take me up on their offer. In this case, I have benefitted them – after all, I saved their life – even though my action was problematic in all sorts of ways.

I have focused so far on cases in which characters are not malicious. When they are malicious, further issues arise. For instance, consider the following case:

RIVAL GANGS. Rivaille and I are in rival mafia organizations. We are both told to kill a certain politician, Paul. I know that if I don't kill Paul, Rivaille will. I kill Paul.

In this case, the worse than nothing account seems to predict that my killing Paul does not harm him; after all, if I had done nothing, Paul would still have ended up dead. But this might, at least initially, appear counter-intuitive; it might appear intuitive that I have harmed Paul.³¹

³¹One might suggest that, in spelling out the case as I have, I have ignored certain possible actions that I might perform. The only options I have considered are killing Paul and doing nothing. But perhaps there are other options available, options which allow Paul to survive. For instance, perhaps there is a third option in which I warn the police of the impending attack and they save his life.

I should start by acknowledging that in many real-life cases such a third option will be present. But not in all cases; sometimes the police won't get there in time, or there is some other reason they cannot help.

Furthermore, I think that if the account handles cases in which there is no such third option, it will also handle cases in which there is a third option. The reason for this rests on a point from my previous subsection: it is important to distinguish harm from failure to benefit. Informing the police about the pending attack would benefit Paul – it will do so by saving him from danger. So if I fail to take this option, I fail to benefit Paul. But

I have three comments. First, it's not clear to me that my account does any worse than my rivals. For example, consider the worse than the alternative account. This says that I harmed Paul if I left him worse off than if I had acted differently. But I don't leave him worse off than if I had acted differently; if I had acted differently (by not killing him), he would have been killed anyways (by Rivaille).

Second, the claim that I harmed Rivaille in this case is difficult to square with intuitions about the other cases I have just presented, such as the case in which I try to save the fragile pedestrian by pushing her out of the way and she dies thanks to my push. In both cases, the character would have died anyways. And it's hard to see how one case would count as harm and the other not.

Third and finally, RIVAL GANGS involves multiple agents. And in such cases, we have more options available for pinpointing which actions cause harm. To develop this point, consider the following:

COLLABORATIVE MURDER. I and my partner are asked by our mafia boss to hunt down and murder a witness. My partner and I prefer to kill as painlessly as possible, and so we tend work as a team. First, I render the witness unconscious using a toxin. Next my partner stabs them to death. Some details: once the toxin is administered, it cannot be reversed. Rather, the way the toxin works is as follows: it first renders someone unconscious for a number of hours. Then, if they are still alive, they wake up in agonizing pain. They stay in this pain for a number of hours and then they die. Further details: if I had chose not to administer the toxin, my partner would have stabbed the witness to death while the witness was fully conscious, which would be considerably more painful.

it doesn't follow that I harm him; harms are different from failures to benefit.

I suppose someone could respond that adding in this third option – the option to benefit – alters whether the second option counts as a harm. In other words, perhaps one thinks in the two option case – I kill Paul or Rivaille does – my killing Paul is not a harm. But perhaps one also thinks in the three option case – I kill Paul, Rivaille kills Paul, or I go to the police – my killing Paul does count as a harm.

To me this new response seems somewhat unmotivated. Once we agree that failure to benefit is different from harm, it's hard to see how adding in an option to benefit could make a difference as to whether a given action counts as a harm. Thanks to an anonymous referee for pressing me on this point.

Such a case seems to have the following features: if I hadn't acted, the witness would have been worse off; they would have been stabbed to death. Likewise, if my partner hadn't acted, the witness would have been worse off; they would have been poisoned and then faced extreme pain. Finally, if neither of us had acted, the witness would have been better off; they would have lived. One last feature worth pointing out: the witness in this case was harmed.

What should the worse than nothing account say about this case? It seems to me that, in order to explain our judgment that the victim was harmed, it should appeal to the fact that if neither of us had done anything, the victim would have lived. More fully, the point is the following: when multiple actions are performed which affect some subject, we can ask the following question: was the set of actions harmful? And to answer this question, the worse than nothing account would advise us to ask: what if none of the actors had done anything to the subject in question; in this case, what would the result would be? If the answer is that the subject would have had more well being, then the result is that the set of actions was harmful. This delivers the correct result that harm was caused to the witness by my poisoning him and then my partner stabbing him.

That said, I should acknowledge that, in some sense, it seems that my giving the witness the toxin benefitted the witness; after all, if I hadn't done so, the witness would have been awake when my partner stabbed them, which would have been excruciating. The worse than nothing account can capture this intuition; after all, if I hadn't performed my action, the witness would have ended up worse than if I'd done nothing.

So far, the worse than nothing account seems to accommodate various intuitions, such as (i) the intuition that the witness was harmed and (ii) the intuition that my giving the witness the toxin, in some sense, benefitted the witness. We also need to accommodate the intuition that my giving the witness the toxin, in some sense, harmed the witness. One thing we can already say to help accommodate this intuition is the following: my giving the witness the toxin was part of a group of actions that together harmed the witness. Nonetheless, a number of further questions remain. When several actions are performed and this set of actions is harmful, does it follow that each of the individual acts is harmful? Or is it just the combined set of actions that is harmful? Also, if each of the actions does count as harmful, how do we apportion the harm? Does each count as equally harmful? Or can some count as more harmful than others?

In sum, then, there are a lot of interesting questions to be raised about cases involving multiple agents. The worse than nothing account appears, at least at first glance, to give attractive verdicts regarding these cases. Nonetheless, much more work is left to be done to fully spell out its implications regarding all such cases.

3 Conclusion

Let me briefly summarize the conclusions of this paper. As I noted early on, harm is an important notion in ethics and the CCA seemed to be the most attractive account of it. But the CCA faces some serious problems, including, notably, the preemption problem. This problem is so serious that it has led some to suggest that the CCA should be abandoned and that perhaps we should abandon the notion of harm altogether. Fortunately, as I have argued, there is a version of the CCA that can solve this problem and, in the meantime, avoid a number of problems that other rival versions of the CCA face. So hope remains for the CCA and the notion of harm in ethics.

References

- [Bennett, 1995] Bennett, J. (1995). *The Act Itself*. Clarendon Press.
- [Boonin, 2014] Boonin, D. (2014). *The Non-Identity Problem and the Ethics of Future People*. Oxford University Press.
- [Bradley, 2012] Bradley, B. (2012). Doing away with harm. *Philosophy and Phenomenological Research*, 85(2):390–412.
- [Dancy, 2004] Dancy, J. (2004). *Ethics Without Principles*. Oxford University Press.
- [Donagan, 1977] Donagan, A. (1977). *The Theory of Morality*. The University of Chicago Press.
- [Feinberg, 1984] Feinberg, J. (1984). *Harm to Others*. Oxford University Press.
- [Feit, 2015] Feit, N. (2015). Plural harm. *Philosophy and Phenomenological Research*, 90(2):361–388.

- [Feit, 2017] Feit, N. (2017). Harming by failing to benefit. *Ethical Theory and Moral Practice*, 22(4):809–823.
- [Feldman, 1991] Feldman, F. (1991). Some puzzles about the evil of death. *The Philosophical Review*, 100(2):205–227.
- [Hanna, 2016] Hanna, N. (2016). Harm: Omission, preemption, freedom. *Philosophy and Phenomenological Research*, 93(2):251–73.
- [Hanser, 2008] Hanser, M. (2008). The metaphysics of harm. *Philosophy and Phenomenological Research*, 77(2):421–450.
- [Howard-Snyder, Summer 2002] Howard-Snyder, F. (Summer 2002). Doing vs. allowing harm. *The Stanford Encyclopedia of Philosophy*.
- [Johansson and Risberg, 2019] Johansson, J. and Risberg, O. (2019). The preemption problem. *Philosophical Studies*, 176:351–365.
- [Kagan, 1989] Kagan, S. (1989). *The Limits of Morality*. Oxford University Press.
- [Klocksiam, 2012] Klocksiam, J. (2012). A defense of the counterfactual comparative account of harm. *American Philosophical Quarterly*, 49(4):285–300.
- [McMahan, 2002] McMahan, J. (2002). *The Ethics of Killing; Problems at the Margins of Life*. Oxford University Press.
- [Norcross, 2005] Norcross, A. (2005). Harming in context. *Philosophical Studies*, 123(1-2):149–73.
- [Northcott, 2015] Northcott, R. (2015). Harm and causation. *Utilitas*, 27(2):147–64.
- [Purves, 2016] Purves, D. (2016). The harms of death. In Cholbi, M., editor, *Immortality and the Philosophy of Death*. Rowman and Littlefield.
- [Purves, 2019] Purves, D. (2019). Harming as making worse off. *Philosophical Studies*, 176(10):2629–2656.
- [Rabenberg, 2015] Rabenberg, M. (2015). Harm. *Journal of Ethics and Social Philosophy*, 8(3):1–32.

- [Roberts, 2007] Roberts, M. A. (2007). The non-identity fallacy: Harm, probability and another look at parfit's depletion example. *Utilitas*, 19(3):267–311.
- [Thomson, 2011] Thomson, J. J. (2011). More on the metaphysics of harm. *Philosophy and Phenomenological Research*, 82(2):436–58.
- [Woollard, 2012] Woollard, F. (2012). Have we solved the non-identity problem. *Ethical Theory and Moral Practice*, 15:677–90.