

Harm: Omission, Preemption, Freedom

NATHAN HANNA

Drexel University

The Counterfactual Comparative Account of Harm says that an event is overall harmful for a subject iff it makes things go worse for the subject, on the whole, than they otherwise would have gone for the subject. Formally:

CCA An event e is overall harmful for a subject S iff S is better off in the closest possible world where e doesn't occur than S is as a result of e .¹

CCA is an attractive account of harm. It has many virtues and many advantages over its competitors. But I won't explore those features here. Ben Bradley gives a nice summary of them, though he ultimately rejects CCA (2012: 396–7, 401). In this paper, I'll reply to two objections to CCA that haven't been adequately addressed in the literature. Some critics think that certain omission and preemption cases raise obviously fatal problems for it. I'll argue that they're wrong.²

First, a caveat: CCA is about *overall* harm, which is different from *pro tanto* harm. I'll explain this distinction later and use it to show that preemption cases don't raise obviously fatal problems for CCA. Until then, my discussion will be confined to overall harm and I'll leave the "overall" implicit.

¹ My informal statement of the view is similar to Bradley's (2012: 396). My formal statement is similar to Klocksiem's (2012: 287). Klocksiem and I use the standard possible worlds analysis of counterfactuals (cf. Lewis 1973; Stalnaker 1968). On this view, a counterfactual $p \square \rightarrow q$ is true iff q is true in the most similar— i.e., "closest"— possible world where p is true (or closest set of worlds where p is true if there's no uniquely closest world). This analysis will help us to more accurately assess the account's implications in the cases that I'll discuss.

² For other objections see Boonin (2014), Bradley (2012), Feit (2015), Gardner (2015), Hanser (2008), Klocksiem (2012), Purves (forthcoming), Shiffrin (1999, 2012), and Tadros (2014). For responses to some of these objections see the works by Boonin, Bradley, Feit, Gardner, Klocksiem, Purves, and Tadros. Klocksiem and Feit offer especially interesting responses to the objections that I'll consider. Though their responses are on the right track, their responses aren't entirely adequate, as we'll see.

1. Omission

Ben Bradley objects to CCA with the following case.

Golf Clubs: Batman buys golf clubs, intending to give them to Robin. That would have made Robin happy. But the Joker convinces Batman to keep the clubs for himself.

Bradley objects to CCA as follows.

[CCA] entails that Batman has harmed Robin, because Robin would have been better off if Batman had not kept the clubs. But it seems implausible to say that Batman has harmed Robin. Merely failing to benefit someone does not constitute harming that person. (2012: 397)

I disagree with Bradley's characterization of CCA's implications. I think he's right that CCA entails that Batman harms Robin, at least if the details of the case are spelled out in the right way (see below). But he's wrong about why.

First, CCA doesn't entail that merely failing to benefit a subject harms the subject. If it did, it would entail that Batman harms Robin for every way he can and doesn't benefit Robin. And it doesn't entail that. As Matthew Hanser points out, it does so only if non-occurrences are events, and they aren't obviously events (2008: 427; cf. Boonin 2014: 53, n2). Even when there are events related to such non-occurrences, CCA doesn't obviously entail that those events are harmful. Suppose that Batman can benefit Robin by giving him the Batmobile. Suppose that he considers doing so and decides not to. CCA doesn't obviously entail that the decision harms Robin. Robin wouldn't obviously have been better off had Batman not decided that. The closest possible world where Batman doesn't decide that isn't obviously one where he gives Robin the Batmobile. It might be one where he never decides, say because he remains indecisive or just forgets to decide. Or it might be one where he doesn't consider giving Robin the Batmobile at all.³

I'll return to these issues later. For now, just notice that Golf Clubs is different from the Batmobile case. Golf Clubs has a special feature: Batman intended to give Robin the clubs and changed his mind. Let's assume for

³ In other words, the following counterfactual might be true: If Batman hadn't decided not to give Robin the Batmobile, it would have been because he didn't consider giving Robin the Batmobile at all. An anonymous referee notes that counterfactuals like this pose certain difficulties. So-called backtracking counterfactuals are claims about how things at earlier times would be different if things at later times were different. David Lewis (1979) says that we typically assume that what happens earlier doesn't depend counterfactually on what happens later. Backtracking counterfactuals come out false on this assumption. However, he notes that in some contexts there are good reasons to abandon the assumption and count certain backtracking counterfactuals as true (1979: 456–7). I'm suggesting that cases like the one above might be like that. The discussion in section 2 will suggest some reasons for thinking this, though I won't explicitly discuss backtracking counterfactuals there.

now that he would have given Robin the clubs had he not changed his mind. If so, CCA entails this:

M Batman's changing his mind harms Robin.

Given our assumption, I think that **M** is true. Here's an argument for it. The Joker's actions harm Robin. They harm him by preventing him from getting a benefit, which they do by getting Batman to change his mind. If Batman's changing his mind doesn't harm Robin, the Joker's actions don't harm Robin either. But since the Joker's actions do harm Robin, Batman's changing his mind does too.

Here's a formal statement of the argument.

- 1) If Batman's changing his mind doesn't harm Robin, the Joker's actions don't harm Robin.
- 2) The Joker's actions harm Robin.
- 3) So, **M** is true: Batman's changing his mind harms Robin.⁴

There are two ways to object to this argument, which I'll consider in turn. You could reject premise 1 and argue that only the Joker's actions harm Robin. Or you could reject premise 2 and deny that the Joker's actions harm Robin.

2. Two Objections to the Argument for **M**

Consider the first objection. If you think that only the Joker's actions harm Robin, you must appeal to a relevant difference between the Joker's actions and Batman's actions to defend your view. Here's one difference that you might appeal to: the Joker's actions interfere with someone else's intentions, but Batman's actions interfere with his own prior intentions. This difference doesn't show that only the Joker's actions harm Robin. Consider why the Joker's actions harm Robin, if they do. If they harm him, it's because of how they affect him: they prevent him from getting a benefit.⁵ Given our assumption that Bat-

⁴ Boonin (2014: 53, n2) gives a different argument for a similar conclusion. I find his argument plausible too.

⁵ Preventing subjects from getting benefits can harm them. Suppose a subject will die without a certain medication, which you're bringing her. I steal it from you. Consequently, she doesn't get it and dies. But she would have lived had I not stolen it. My stealing the medication harms her (cf. Hanser 2008: 427–8; Shiffrin 2012: 368). In light of this example, you might reply that preventing subjects from getting only certain kinds of benefits is harmful for them. You might say that preventing a subject's life from being saved harms the subject but preventing Robin from getting the golf clubs doesn't harm him. I'll address this objection when discussing the second objection to my argument for **M**.

man would have given Robin the clubs had he not changed his mind, Batman's changing his mind affects Robin in the same way. Indeed, the Joker's actions have that effect *because* Batman's changing his mind has that effect. And it's the effect that makes these events harmful for Robin if they're harmful for him. Since both the Joker's actions and Batman's actions affect Robin in this way, both harm him if either does. So appealing to the above difference doesn't vindicate the first objection.

Here's another difference that an objector might appeal to. The Joker actively prevents Robin from getting a benefit, but Batman only passively refrains from benefitting Robin. Justin Klockslem (2012) actually defends CCA against the omission objection by appealing to this difference and arguing that CCA doesn't entail that agents like Batman harm agents like Robin. If he's right, my argument for M fails. So, let's consider what he says about the active/passive distinction.

Klockslem uses a case like Golf Clubs to explain his take on the distinction's significance. His case is simpler, so I'll use it to explain his view. Then I'll discuss the implications that his discussion has for Golf Clubs. Here's his case.

Hat: Archie buys a hat for Jughead, but then decides to keep it for himself.

Klockslem grants that getting the Hat would make Jughead better off (let's assume that getting it would make Jughead happy). But Klockslem thinks that Archie's changing his mind doesn't harm Jughead. Here's why.

[W]e tend to view structurally identical cases differently, depending on whether the agent was active or passive, because this affects our intuitions about what would have happened had things been otherwise. So in [Hat], we assume that Archie was going to refrain from giving the hat to Jughead, conclude that Archie did not affect Jughead's welfare, and classify it as a mere failure to benefit. (2012: 294)

The thought here seems to be this: the fact that Archie changed his mind suggests that his intention was particularly unstable. His reasons for changing his mind suggest that he was unlikely to follow through on his intention and give Jughead the hat. So the closest world where he doesn't change his mind is probably still one where he doesn't give Jughead the hat. Klockslem thinks that these observations help to illustrate a more general point.

In general, whether an event constitutes a harm or a mere failure to benefit depends heavily on contextual details that determine whether or not the relevant possible world in which S is better off is very near the [actual] world. If the world in which S is better off is very similar to the [actual] world, and so an intervention was required to prevent that world from coming about, we are more likely to regard [such an intervention] as a

harm. If the world in which *S* is better off is not very similar to the [actual] world, and so an intervention was required to bring it about, we are likely to regard [omitting to so intervene] as a [mere] failure to benefit. (2012: 294–5)

Klocksziem's take on all of this is ingenious, but not fully satisfactory. I'll explain.

Note that we're not told what Archie's reasons for changing his mind are. If they're predominantly selfish, it's plausible to say that Archie's intention was unstable. Such reasons seem to reveal deeper facts about Archie that make it unlikely that he'd follow through on his intention and give Jughead the hat. I agree with Klocksziem about such cases: CCA doesn't obviously entail that Archie harms Jughead in such cases. This is because the closest possible world where Archie doesn't change his mind isn't obviously one where he gives Jughead the hat. But things can change if Archie's reasons are different. Suppose that Archie meets a pretentious fashion critic who mocks the hat. Archie then worries that Jughead won't like it and decides not to give it to him. Here, Archie's reasons don't suggest that his intention was particularly unstable. The closest world where Archie doesn't decide to keep the hat isn't obviously one where he gives Jughead the hat. It might be one where Archie just isn't swayed by the critic and gives Jughead the hat anyway. If it is, CCA entails that Archie's changing his mind harms Jughead.⁶

The same goes for Golf Clubs, which is structurally similar to the fashion critic variant of Hat. The closest world where Batman doesn't change his mind isn't obviously one where Robin doesn't get the clubs. What world is closest depends on details like Batman's reasons for changing his mind. Suppose that the Joker does something like what the fashion critic does. Batman's reasons for changing his mind in such a case don't suggest that his intention is especially unstable. The closest world where he doesn't change his mind may very well be one where he just isn't convinced by the Joker and gives Robin the clubs anyway. If so, CCA entails *M*. And I think that's the right result, for the reasons I've given.⁷

I've considered some important differences between the Joker's actions and Batman's actions. These differences don't show that only the Joker's

⁶ In such a case, CCA also entails that Archie's meeting the fashion critic and the fashion critic's mocking the hat harm Jughead. Given such results, Hanser objects that CCA excessively multiplies harm (2008: 433–4). See Klocksziem (2012: 293) and Bradley (2012: 405) for good responses.

⁷ Notice: if we stipulate that the Joker will ensure that Robin doesn't get the clubs no matter what Batman does, then CCA doesn't entail *M*. If you're worried that CCA entails that the Joker's actions don't harm Robin in this version of the case (and you think that they do), see my discussion of the preemption objection below, in sections 4 and 5. My response to the preemption objection deals with this worry.

actions harm Robin. Since they seem like the most important differences, the first objection to my argument for M seems to fail. So I'll move on to the second objection.

Recall that the second premise of my argument for M says that the Joker's actions harm Robin. In defense of this, I said that the Joker's actions harm Robin because they prevent him from getting a benefit. The second objection to my argument for M says that the Joker's actions don't harm Robin. Klocksiem seems sympathetic to this objection and presses it in the obvious way. He considers a similar case where an event prevents someone from getting a trivial benefit. He says that the event doesn't "appear to rise to the level of genuine harm" (2012: 294). One might say the same about the Joker's actions.

If you think this, I disagree with you. I think that the Joker's actions harm Robin, but just a little bit (cf. Boonin 2014: 53, n2). CCA's fate doesn't turn on whether I'm right about this, though. If you think that the Joker's actions don't harm Robin, this gives you a reason to reject M. But it doesn't give you a reason to reject CCA. There are at least two ways to reconcile CCA with the view that the Joker's actions don't harm Robin. One: endorse an account of well being that entails that getting the golf clubs wouldn't have made a difference to Robin's well being (cf. Tadros 2014: 178).⁸ Two: modify CCA. An obvious way to do that is to say that an event is overall harmful for a subject iff it makes things go worse *by at least a certain amount* for the subject than they otherwise would have gone for the subject.⁹ These strategies are compatible. Both can be used to preclude CCA from entailing claims like M in cases like Golf Clubs and Hat. To some extent, then, they can be used to disarm objections like Bradley's. No plausible application of these strategies will completely eliminate such entailments, though. In variants of the cases involving sufficiently large benefits, CCA will entail claims like M. And that seems like the right result. Arguments like my argument for M will be harder to reject in such cases.

To sum up the discussion so far, there are good reasons to accept that CCA sometimes entails claims like M and that claims like M can be true. If so, the omission objection to CCA isn't obviously fatal. Intuitions that are at odds with claims like M are doubtful. I suspect that many who find such

⁸ Klocksiem is willing to grant that getting the clubs would make Robin better off. He grants that receiving the trivial benefit in his case would make the subject better off and that getting the hat would make Jughead better off (2012: 286–7, 293). This concession is incompatible with the above strategy.

⁹ This modification is a simple quantitative one. Qualitative modifications are also possible. Example: an event is overall harmful for a subject iff it makes things go worse *in certain ways* for the subject than they otherwise would have gone for the subject. I doubt that there are any good reasons to accept such a view. But I won't explore this. What I say about the quantitative principle applies to lots of qualitative ones too.

claims counterintuitive are led astray at least partly by our tendency to moralize harm. This is our tendency to be more likely to call acts harmful if we think that they're wrong or motivated by bad intentions.¹⁰ At least, this is a plausible explanation of why one might mistakenly think that actions like the Joker's are harmful and actions like Batman's aren't, even when they have the same effect on a subject's welfare.¹¹

My reply to the omission objection invites some further objections, though. I'll consider them in the next section. Then I'll deal with the preemption objection.

3. Seriousness and Moral Significance

Maybe I've missed something. Maybe the fact that CCA sometimes entails claims like M is symptomatic of a fatal problem I've overlooked. Consider a variant of Golf Clubs. Suppose Batman intends to give Robin lots of money and the Joker convinces Batman to change his mind. Assume that Batman would have given Robin the money had he not changed his mind. One might object that if the amount of money is large enough, CCA entails that Batman's changing his mind *seriously* harms Robin. One might object that this is a fatal problem for CCA.

It's not. The claim that the harm is serious can be understood in two ways. First, it might be a claim about the harm's comparative size. It might mean that the harm is much bigger than in cases where the benefit that Robin is prevented from getting is much smaller. But that's unproblematic. *If* Robin is harmed here, he's harmed much more than he is in cases where the benefit that he's prevented from getting is much smaller. This isn't to say that the harm is especially significant in ways that would make this result problematic, though. Considering a second interpretation of the seriousness claim makes this clear.

¹⁰ Klockslem falls prey to this tendency. He argues that acts like Archie's changing his mind aren't harmful because it would be supererogatory to give the subject the relevant benefit (2012: 294).

¹¹ Bradley agrees that moralizing leads our intuitions about harm astray (2012: 410; Norcross 2005: 151). And he argues that we shouldn't moralize harm (2012: 394–5; cf. Boonin 2014: 61–5). Moralized accounts fail to count lots of harmful events as harmful. Counterexamples include non-act events like explosions and earthquakes. They also include permissible acts done for good reasons, e.g., justifiable uses of lethal force in self-defense. Comparisons like the following also speak against moralizing harm. “[I]f a boulder falls on someone, the person is harmed by the boulder's falling in just the same way that she would have been harmed by someone's pushing the boulder on her” (Bradley 2012: 408 n23). Bradley says that the permissibility of the act and the quality of the agent's intentions don't affect whether or how much the subject has been harmed. They only affect how blameworthy the agent is. I agree. Thanks to an anonymous referee for pressing me here.

Second, the claim that the harm is serious could be a claim about the harm's moral significance. It might mean that the harm is especially morally significant, so that there's a very strong, possibly decisive reason for Batman to not change his mind. I agree that that claim is false.¹² But CCA doesn't entail it. CCA is a claim about what makes events harmful. It's not a claim about the moral significance of harms or about the strength of reasons not to harm.

Given this, one might object that CCA doesn't tell us enough about harm. Some critics object to CCA on such grounds. Matthew Hanser objects to a version of CCA on the grounds that a good account of harm should do more than identify harms. He says that it should be able to classify harms into distinctive types of varying moral significance (2008: 428). And Seana Shiffrin objects that CCA doesn't capture harm's apparent moral priority: the apparent fact that, other things equal, our reasons not to harm seem stronger than our reasons to benefit (2012: 361, 371–2; cf. Harman 2004: 98).

My reply is that it's not obviously a problem for CCA if it can't tell us some important things about harm. Maybe the right theory of harm can only tell us certain things that we want to know when conjoined with theories about other things, e.g., well-being, desert, and rights (cf. Bradley 2012: 410). We might need such theories before we can use CCA to determine the moral significance of particular harms and related facts about the strength of reasons not to impose them. I see no reason to assume that CCA should be able to tell us such things in isolation from a larger body of moral theory. Obviously, I don't have space to elaborate on those other parts of moral theory here.

These remarks invite another objection, though. One might object that harm is especially morally important and that CCA trivializes harm, or at least that my defense of CCA does. Bradley and Kloocksiem say that popular principles like Mill's Harm Principle suggest that harm is especially important, such that we have strong, perhaps overriding reasons to prevent harm and to avoid harming (Bradley 2012: 390–1, Kloocksiem 2012: 285; cf. Boonin 2014: 84, Harman 2004: 93–4, 2009: 139).¹³ Call this the Strong View. I think that CCA's implications in some cases are inconsistent with this view. That's not obviously a problem for CCA, though. Maybe the Strong View is false (cf. Tadros 2014: 176). Maybe some harms aren't as important as that view predicts. It shouldn't be surprising if there are lots of misconceptions about harm. For one thing, the worst kinds of harm can be the most salient when we're thinking about harm. This can make us overlook other kinds of harm and lead our intuitions about harm astray. For another, philosophers haven't paid much attention to the nat-

¹² Things are different when it comes to the Joker. It seems that the Joker shouldn't interfere, at least not for the reasons that he would characteristically have for interfering. Not so for Batman. There are reasons for him not to change his mind. But it seems that they can be outweighed by other reasons that permit him not to make the sacrifice.

¹³ Bradley also appeals to the Precautionary Principle and a version of the Hippocratic Oath. Kloocksiem also appeals to W.D. Ross's statement of the duty of non-maleficence.

ure of harm. Bradley motivates his entire discussion by emphasizing this: “Despite the importance harm is supposed to have, almost nobody bothers to say what it is” (2012: 391).

These observations counsel caution. We shouldn’t quickly reject CCA on the basis of unexamined assumptions about harm, even if they seem to be widely shared and seem to motivate popular principles about harm. Consider the principles that Bradley and Klocksiem appeal to. I think that many popular principles about harm are oversimplified and overstated, persistently misinterpreted, or implicitly restricted to certain kinds of harm. Take Mill’s Harm Principle.

The principle requires liberty. . . of doing as we like, subject to such consequences as may follow: without impediment from our fellow creatures, so long as what we do does not harm them. (Mill 1859: 265)

Many philosophers think that this principle assumes something like the Strong View. Insofar as one finds the principle and this interpretation of it attractive, one will be attracted to that view. But many philosophers take issue with the principle or with common interpretations of it. To take just two examples, David Brink says that the principle is overly simplistic (2014: S3.6). And Piers Norris Turner argues that it’s persistently misinterpreted. He argues that the principle relies on an “expansive” conception of harm that’s consistent with my position (2014: 300). In light of these controversies, the principle doesn’t obviously give us good reason to think that CCA or my defense of it misrepresents harm’s moral importance. I think that it’s an open question how important harm is, generally speaking. CCA might help us answer that question.

That’s all I’ll say about omissions. To sum up, omission cases don’t pose obviously fatal problems for CCA. CCA doesn’t entail the unacceptable claim that one harms merely by failing to benefit. Its implications are narrower. In at least some cases, CCA entails claims like M. These implications might seem counterintuitive. But there are good reasons to accept them and to reject contrary intuitions.

I’ll now consider the preemption objection to CCA.

4. Preemption

Bradley also objects to CCA with the following case.

Heart Attack: Batman dies of a heart attack just before he would have been killed by a flaming cannonball.

Bradley objects to CCA as follows.

[CCA] entails that the heart attack was not harmful to Batman. It didn’t make things go worse for him. But intuitively, the heart attack was

harmful. The fact that he would have been harmed by the flaming cannonball anyway does not seem relevant to whether the heart attack was actually harmful. (2012: 397)

Among others, Molly Gardner (forthcoming: 9), Matthew Hanser (2008: 434–7), Frances Kamm (2014: 504), Michael Rabenberg (2015: 11), and Seana Shiffrin (2012: 367–8) agree that CCA gets things wrong in preemption cases.¹⁴ I disagree.

With one qualification that I'll state later, I endorse Klocksiesm's response to this sort of objection. Klocksiesm says that events like the heart attack are *pro tanto* harmful, not overall harmful (2012: 295–6). Bradley suggests one way to understand this distinction (2012: 393–4; cf. Klocksiesm 2012: 289–90; Norcross 2005: 150–1; Purves forthcoming: 18).

- An event *e* is *pro tanto* harmful for a subject *S* iff *e* has an overall-harm-making feature for *S*.
- An event *e* is overall harmful for a subject *S* iff *e*'s *pro tanto* harms to *S* outweigh its *pro tanto* benefits to *S*.

I'll explain this take on the distinction in more detail below. But first, consider an example that will give us an intuitive grip on the distinction (cf. Harman 2009: 139). Suppose *S* has invasive life-saving surgery. The surgery is *pro tanto* harmful for *S* because it's invasive. But it's overall beneficial for *S* because it's life-saving. This example shows that *pro tanto* harm is a genuine kind of harm and that the concept of *pro tanto* harm is theoretically useful. The invasiveness is a *pro tanto* harmful feature of the surgery. It's a *pro tanto* reason against performing the surgery and a *pro tanto* reason to do things like develop and use comparably effective non-invasive treatments. This reason can make performing the surgery wrong when such alternatives are available or when the surgery doesn't have outweighing benefits for the patient (cf. Boonin 2014: 84).

Now, consider Bradley's suggested take on the distinction. It needs some unpacking. Consider the claim that an event is *pro tanto* harmful for *S* iff it has an overall-harm-making feature for *S*. What's an overall-harm-making

¹⁴ They discuss different kinds of cases than Bradley's. In their cases, an agent's free act preempts another agent's free act. Their cases are also different from Norcross's case, discussed below. In that case an agent acts in a way that preempts another act that he would otherwise perform. As we'll see, a version of Norcross's case raises a serious problem for CCA. My solution avoids it. But my solution gives the same verdict in Gardner's, Hanser's, Kamm's, Rabenberg's, and Shiffrin's cases that CCA gives in Heart Attack. I think this is the right verdict for the same reason that it's the right verdict in Heart Attack. See the discussion below and in section 5.

feature for S?¹⁵ I endorse a variant of Bradley's view from his earlier work. Here's his view (2009: 66).

PTH An event e is pro tanto harmful for a subject S iff:

- 1) e is intrinsically bad for S , or
- 2) e brings about something intrinsically bad for S , or
- 3) e prevents something intrinsically good for S .

PTH is a list of overall-harm-making features for S . Theories of well-being will tell us what things are intrinsically good or bad for S . I said that I endorse a variant of PTH because it seems like PTH can't adequately deal with certain overdetermination and preemption cases. Here's an overdetermination case.

S is shot twice simultaneously. S is killed instantly and painlessly. Each shot alone would have had this effect.

Since the shots kill S , each shot seems pro tanto harmful for S . But PTH doesn't obviously entail this.¹⁶ When death or events that cause it are pro tanto harmful for a subject, it's at least partly because they satisfy PTH's condition 3, i.e., because they deprive the subject of future benefits. But neither shot does this. Each is such that, had it not occurred, S would have been killed anyway. So, PTH seems to go wrong in this case. Preemption cases like the following raise the same problem.

Event e_1 kills subject S instantly and painlessly. Another event e_2 would otherwise have had the same effect moments later. S would have lived had neither event occurred.

e_1 seems pro tanto harmful for S . But again, PTH doesn't obviously entail this. Something similar applies to Heart Attack. Batman's heart attack satisfies condition 2, assuming that it's painful. But that's not the only reason why it's pro tanto harmful for him. It also seems pro tanto harmful for him because it kills him.

I suspect that the most promising way to deal with these problems for PTH is with an account of plural harm. Such an account says when multiple events *together* are overall harmful for a subject. Neil Feit (2015) defends a

¹⁵ Thanks to an anonymous referee for pressing me here.

¹⁶ PTH will entail this if we combine it with a theory of well-being on which mere bodily damage is intrinsically bad for a subject. I won't assume such a theory, though. Even if such a theory is true, PTH still can't account for all the ways in which the shots are pro tanto harmful for S . See the next sentence.

sophisticated and attractive account. I won't discuss its details or try to defend it here, though.¹⁷ I mention it only because such accounts suggest that we should add another condition to PTH. We should also say that an event *e* is pro tanto harmful for a subject *S* if something like the following is true.

- 4) *e* is one among a group of events that together satisfy at least one of PTH's other conditions.

With this addition, we can say that the heart attack is pro tanto harmful for Batman because it's one among a group of events that together deprive him of future benefits. The heart attack is part of this group because it kills Batman. The events involving the cannonball are part of this group because they would have killed him had the heart attack not done so.

So much for pro tanto harm. Next, consider the claim that an event *e* is overall harmful for a subject *S* iff its pro tanto harms to *S* outweigh its pro tanto benefits to *S*. This doesn't make clear how pro tanto harms can outweigh pro tanto benefits. So here's a clearer, if cumbersome way to put it.¹⁸

An event *e* is overall harmful for a subject *S* iff the sum of *e*'s intrinsic badness for *S*, the intrinsic badness for *S* of what *e* brings about, and the intrinsic goodness for *S* of what *e* prevents outweighs the sum of *e*'s intrinsic goodness for *S*, the intrinsic goodness for *S* of what *e* brings about, and the intrinsic badness for *S* of what *e* prevents.

Now, let's reconsider Heart Attack in light of the above discussion. I said that, subject to a qualification, the heart attack is pro tanto harmful for Batman but not overall harmful for him. PTH and the above take on overall harm explain why CCA entails this and why this is a plausible result.

The heart attack is pro tanto harmful for Batman because it satisfies PTH's conditions 2 and 4. It satisfies 2 because it causes him pain. It satisfies 4 because it's one among a group of events that together are overall harmful for him. The heart attack itself isn't overall harmful for him, though. It's not obviously intrinsically bad for him. Nor does it prevent anything intrinsically good for him, since he would have died moments later anyway. And though it causes him pain, which is intrinsically bad for him, it presumably also prevents the cannonball from doing so. This is where the

¹⁷ One of the most important details of such an account is how it identifies the events that together harm subjects. A good account will avoid being overinclusive or underinclusive. That is, it won't include too many or too few events as members of the harmful group. Feit's proposal is attractive in this respect (2015: 371, 374–9).

¹⁸ Thanks to an anonymous referee for pressing me here.

qualification that I mentioned comes in. Whether the heart attack is overall harmful for Batman depends on whether it causes him more or less pain than it prevents. If, as seems plausible, it causes him as much or less pain than the cannonball would have, then it's not overall harmful for him and may even be overall beneficial for him. However, if it preempts a less painful death, then it is overall harmful for him.¹⁹

These results are plausible. The claim that the heart attack is pro tanto harmful for Batman accounts for our intuitions about the case. It accounts for our intuition that the heart attack harms Batman, since pro tanto harm is a genuine kind of harm. And it accounts for the intuition that there's something about the case that makes the heart attack different. Since Batman would have been killed anyway, the heart attack doesn't seem to matter in the same way that it does in a case where he otherwise would have lived. Bradley explains this by appealing to the claim that harm is prudentially important (2012: 406). Suppose Batman knew that he was going to be killed either way. Would he care if he died of a heart attack instead of a cannonball strike? Bradley says that he probably wouldn't and shouldn't care.²⁰ Bradley grants that this supports the claim that the heart attack isn't overall harmful.

In light of this, Bradley grants that my response to the preemption objection has some plausibility. Despite this, he thinks that the response fails. To show this, he appeals to a case discussed by Alastair Norcross (cf. Norcross 2005: 165–66).

Dark 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.

CCA entails that Knight's choking Phil is overall beneficial for Phil, assuming that we treat benefit and harm analogously (I'll assume this throughout). Here's Bradley again.

This case seems to show that according to [CCA], one can make one's pain-causing actions overall beneficial rather than harmful merely by

¹⁹ If things other than pain are also intrinsically bad for subjects and if such things are implicated in this case, the above points can be put in terms of those other things too. Bodily damage might be intrinsically bad for subjects, say. Harman might be working with such a view (2004: 91–2, 97, 2009: 139; cf. Boonin 2014: 75).

²⁰ This must be qualified in light of my remarks above. Batman should and probably would care about whether the heart attack kills him in a more or less painful way. But he shouldn't and probably wouldn't care about the mere fact that it kills him, since he'd otherwise die moments later anyway.

becoming a rage-filled lunatic who would otherwise have done much worse things. This is hard to square with our ordinary practice of harm attribution... [It] seems that there is serious overall harm here. (2012: 407)

Klocksien doesn't discuss this case. Feit does and he thinks that his account of plural harm can deal with it. But his account entails that the choking is overall beneficial for Batman rather than overall harmful for him (2015: 381; cf. Boonin 2014: 62–3). Norcross offers a solution to the problem that the case raises for CCA, but he says that his solution makes the concept of harm useless for moral theory.²¹ I think that there's a better way to deal with the case. My solution doesn't make harm a useless concept and it doesn't entail that the choking is overall beneficial for Batman. I'll present it in the next section.

5. Freedom

To deal with Dark Knight, we must clarify some of its details first. Bradley likens Knight to a rage-filled lunatic. Norcross frames things in term of Knight's attempts to control his behavior. These characterizations invite confusion about something important: how much control Knight has over his behavior. Suppose that he has very little control. He's overcome by rage and the only way that he can prevent himself from dismembering Phil is by choking Phil. If so, CCA's verdict seems right. Knight's choking Phil is overall beneficial for Phil.²² Dark Knight is only a problem for CCA if Knight

²¹ Here's his view: Act A harms subject S iff S is worse off as a result of A than S would have been if *the appropriate alternative* to A had been performed instead (2005: 169). Norcross says that conversational context determines what the appropriate alternative is by making some alternatives especially salient. But he doesn't spell out the details of this view (2005: 171). Whatever the details, he thinks that on this view "harm and benefit do not have the kind of metaphysical grounding required to play fundamental roles in ethical theory, nor do judgments of harm and benefit make any distinctive contribution to reasons for action" (2005: 171). Bradley calls this a contrastive account and says that it fails to make harm normatively important (2012: 407–8, 410).

²² If you're skeptical, consider another case. You anger me and I start turning into the Hulk. If I do nothing, I'll transform, lose control, and smash you. The only way that I can prevent this is by locking you in a nearby basement where you'll be out of sight. So I do that, roughing you up in the process because you ignore my warnings and resist. CCA entails that my action is overall beneficial for you. And that verdict seems right. There are no relevant differences between this case and the out of control version of the Knight case that support different verdicts.

has more control over his behavior. So I'll stipulate that Knight would freely do whatever he does: he freely chokes Phil and would have freely dismembered Phil otherwise.²³ Given this stipulation, CCA's verdict is highly counterintuitive.

We can avoid this result by using a different counterfactual comparison than the one that CCA uses. We should compare how well off Phil is after Knight chokes him to how well off he would have been had Knight not tried to hurt him at all. There are two reasons to use this comparison. First, choking Phil and dismembering Phil are different ways for Knight to do the same thing for the same reasons. They're both attempts by Knight to hurt Phil (see my discussion of intentions below for more on this). Second, Knight's actual and counterfactual acts are free. Because of this we assume that, at the point of decision, he could have refrained from trying to hurt Phil. This follows from an assumption that we make about free will: that when agents act freely, they could have acted otherwise. Since Knight acted freely and could have refrained from trying to hurt Phil, we should use a comparison where he so refrains. That gets the intuitively right result. And it does so in an intuitively attractive way. To illustrate, the reasons that I've appealed to are reasons why the choking is prudentially important for Phil. Phil has good reasons to prefer not to have been choked by Knight, despite what Knight would otherwise have done. This is because Knight could have refrained from trying to hurt him and because Phil would have been better off had Knight so refrained.

These considerations suggest a way to revise the counterfactual comparative account. Since the reasons that I've given for using a different comparison in *Dark Knight* don't apply to mere events, we should say that CCA only applies to mere events and that a different principle applies to free acts.²⁴ I propose this:

²³ Some philosophers think that counterfactuals of freedom like this have no truth-value. I'm sympathetic to this view, but won't discuss it. Even if it's right, counterfactuals about what agents would *probably* freely do could be true. My discussion can be framed in such terms. So framed, the counterfactual comparative account's implications are more complicated. In cases involving free acts, it will entail that certain acts and events are *probably* harmful or beneficial.

²⁴ Bradley says that accounts of harm should be ontologically neutral. Initially, he says that accounts are ontologically neutral if they count both acts and non-act events as harmful and if they recognize that different sorts of beings can be harmed, not just typical human adults (2012: 394–5). I agree that accounts of harm should be neutral in these ways. My view is. But Bradley later says that accounts that treat acts and non-act events differently aren't ontologically neutral (2012: 408). I don't think that accounts of harm should be neutral in this way. The reasons that I've given above are good reasons to treat acts and non-acts differently.

CCA+ 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.²⁵

Call the conjunction of CCA and CCA+ the *Revised Counterfactual Comparative Account of Harm* or simply the *revised account*. I'll discuss how to apply it in a moment, focusing specifically on how to classify free acts as acts of certain types. Then, I'll argue that it gets the right results in various cases.

But first I need to make a qualification. When I gave reasons to use a different comparison in *Dark Knight*, I appealed to an assumption: that when agents act freely, they could have acted otherwise. This is a controversial claim about free will. I think that it's true. But I won't defend it. If you reject it, take me to be trying to show how those who accept it should revise the counterfactual comparative account to deal with *Dark Knight*. If the assumption is false, I'd say that no revisions are needed because CCA gets the right result in *Dark Knight*. I'd say that Bradley's objection fails because it appeals to mistaken intuitions produced by the assumption.

Now I'll discuss how to apply the revised account. An obvious objection to it says that there's no good way to classify free acts for CCA+'s purposes. Hanser levels such an objection against a similar principle meant to accommodate his intuition that events like Batman's heart attack are overall harmful.²⁶ Hanser doubts that there's a way to classify events as relevantly similar that doesn't presuppose an account of harm. You might have similar suspicions about CCA+ and about my claim that choking Phil and dismembering Phil are different ways for Knight to do the same thing.

The suspicions are unfounded. We can classify free acts in the needed way without presupposing an account of harm. We can do this by looking at the agent's intentions. When Knight chokes Phil, he acts on various intentions such as an intention to hurt Phil. So here's an initial proposal.

²⁵ I don't really have to formulate a different principle. Alternatively, I could modify CCA as follows: An event *e* is overall harmful for a subject S iff S is better off in the closest *relevant* possible world where *e* doesn't occur than S is as a result of *e*. I could then say that the relevance criteria for mere events are different than those for free acts (and spell out the latter by appealing to intentions; see below). I prefer putting things in terms of CCA+ for ease of exposition, though—and because CCA+ is based on an assumption about free will that some will reject. If you reject the assumption, see below. If the assumption is true, we'll also need variants of PTH and of my explanation of Bradley's take on overall harm. These variants will have modified prevention clauses. They'll determine what an agent's free act prevents by looking at the closest possible world where the agent doesn't freely perform an act of that type.

²⁶ Here's the principle: An event *e* is overall harmful for a subject S iff *e* makes S worse off than S would have been had neither *e* nor any relevantly similar event occurred (Hanser 2008: 435).

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

Suppose that Knight choked Phil because Knight intended to hurt him. Phil would have been better off had Knight not tried to hurt him. So on the initial proposal the revised account entails that Knight's choking Phil overall harms Phil.

Unfortunately, the initial proposal doesn't work. Knight intended to hurt Phil. But he also intended to restrain himself. He choked Phil instead of dismembering Phil. So his act counts as an act of restraint on the proposal. Had Knight not restrained himself, Phil would have been worse off because Knight would have dismembered him. So on this proposal the revised account still entails the counterintuitive claim that Knight's choking Phil overall benefits Phil. Even worse, on this proposal the revised account yields contradictory implications. Knight's choking Phil is overall beneficial for Phil because it's an act of restraint. And it's overall harmful for him because it's an attempt to hurt him.

Fortunately, we can improve the initial proposal. To classify the choking, we must account for some important details about Knight's intentions. Some of the intentions that he acts on are more fundamental than others in the sense that they have a motivational priority that the others don't have. Knight's intention to hurt Phil is more fundamental in this sense than his intention to restrain himself. One way to see this is by trying to explain Knight's actual and counterfactual behavior. Why does he choke Phil? Why would he otherwise have dismembered Phil? And why does he restrain himself in the way that he does, i.e., not so much that he doesn't try to hurt Phil at all? Answer: because he intends to hurt Phil and because this intention is fundamental in the above sense. We can't give a fully satisfactory explanation of Knight's actual and counterfactual behavior without acknowledging this fact. This suggests a way to improve the initial proposal about how to classify free acts. We can avoid the problems with that proposal by modifying it as follows.

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.

This proposal classifies Knight's choking Phil as an attempt to hurt Phil, not as an act of restraint. On this proposal the revised account gets the right result: Knight's choking Phil is overall harmful for Phil.

This improved proposal coheres well with our intuitions about how best to describe Knight's act for the purpose of harm assessment. For that purpose, it's inaccurate and disingenuous to simply call Knight's choking Phil

an act of restraint. Doing that omits crucial details about the nature of the restraint and about the overall character of Knight's intentions. Classifying Knight's act as an attempt to hurt Phil doesn't do that. On this improved proposal, the revised account is a promising way to deal with the problem that Dark Knight raises for CCA. At least, it's promising given the revised account's assumption about free will.

Moreover, the account seems to have plausible results in the other cases that I've discussed.²⁷ For example, CCA+ seems to have the same results in cases like Golf Clubs and Hat that I defended earlier. That said, you might be skeptical about the account's ability to deal with other cases. Recall the case where Batman considers giving Robin the Batmobile and decides not to. I argued that CCA doesn't obviously entail that Batman's decision harms Robin. But you might worry that CCA+ has unacceptable results in this case.²⁸ Suppose that Batman's fundamental intention is to keep the Batmobile. You might object that CCA+ entails that Batman's decision harms Robin because Robin would have gotten the Batmobile had Batman not acted with that intention. This objection fails, though. My earlier defense still works. The closest world where Batman doesn't act on an intention to keep the Batmobile isn't obviously one where he acts on an intention to give Robin the Batmobile. As before, the closest such world may be one where Batman remains indecisive, forgets to decide, or never consider giving Robin the Batmobile at all.

In light of all this, the revised account seems promising. However, there are some especially challenging objections to it. I'll discuss two in the next section.

6. Two Objections to the Revised Account

The first objection to the revised account focuses on CCA+'s appeal to fundamental intentions. Consider the following case.²⁹

Dark Knight 2: Knight wants to steal Phil's life savings. Knight knows that the only way he would refrain from freely doing so is by kicking Phil. If he does that, he'll feel sorry for Phil and refrain from freely stealing Phil's savings. So Knight freely kicks Phil with the fundamental intention of preventing himself from freely stealing Phil's savings.

²⁷ This is one reason to think that the addition of CCA+ isn't just an ad hoc way of dealing with Dark Knight. The reasons that I gave for treating free acts differently are also reasons to think that the account isn't ad hoc. Thanks to an anonymous referee for raising this worry.

²⁸ Thanks to an audience member at the Central States Philosophical Association for this worry.

²⁹ Thanks to Eric Vogelstein for this objection.

Had Knight not done that, he would have freely stolen Phil's savings with the fundamental intention of enriching himself.

The objection claims that the kicking and the stealing can't be classified as acts of the same type, since the fundamental intentions for them are different. Because of this, the objection concludes, the revised account entails the counterintuitive claim that the kicking is overall beneficial and merely pro tanto harmful for Phil.

The crucial detail here is that kicking Phil is the *only way* that Knight would refrain from stealing from Phil. But why is this the only way? There seem to be two possibilities here. One is that Knight lacks a significant degree of control *now* over what he would freely do *later*—hence the talk of preventing himself from stealing Phil's savings. If so, CCA+ does entail that Knight's kicking Phil is overall beneficial and merely pro tanto harmful for Phil. But that's the right result. If the case involves a significant lack of control, it's like a case where Knight kicks Phil because that's the only way to prevent someone else from freely stealing Phil's savings. Kicking Phil is overall beneficial and merely pro tanto harmful for Phil in such a case.

Alternatively, the *only way* claim might be true because Knight is *willing* to refrain from stealing from Phil only if he kicks Phil. If so, the objection's description of Knight's intentions is wrong. Knight's intention to prevent himself from stealing Phil's savings isn't fundamental—much like his intention to restrain himself in Norcross's case. His fundamental intention is actually something like this: *to either steal from Phil or kick Phil* (or maybe it's like this: *to either enrich himself at Phil's expense or hurt Phil*; the basic point is that fundamental intentions can be more complex than the objection assumes). CCA+ entails that Knight harms Phil overall when he freely performs an act of that type, i.e., a *stealing-or-kicking*. Had Knight not freely performed an act of that type, he wouldn't have freely kicked Phil or freely stolen from Phil. And Phil would have been better off.

So here's my general response to problem cases like Dark Knight 2. In such cases, there are two possibilities. One: the agent's fundamental intentions for her actual and counterfactual free acts are really the same. They're just more complex than the objection assumes. Two: the agent lacks control in an important way. I suspect that a genuine divergence of fundamental intentions is symptomatic of such a lack of control. On the first possibility, the revised account entails that the agent's free act is overall harmful for the subject. On the second, the account entails that the agent's free act is overall beneficial and merely pro tanto harmful for the subject. Either way, the account gets the right results.

The second objection to the revised account exploits the fact that the account treats free acts and mere events differently.³⁰ Consider another case.

Dark Knight 3: Knight gets mad at Phil and wants to hurt him. Knight throws a chair at him, causing the chair to fly towards him and hit him. Knight would have freely dismembered Phil had Knight not thrown a chair at him.

CCA+ applies to Knight's free act and counts it as overall harmful. But CCA applies to the mere events of the chair flying towards Phil and the chair hitting Phil. The objection says that CCA counts these mere events as overall beneficial and merely pro tanto harmful for Phil because Knight would have dismembered Phil had they not occurred. The objection claims that this result is unacceptable because the events are overall harmful for Phil.

Again, there are two ways of filling in the details here. On the first, the revised account doesn't have the implications that the objection says it does. On the second, it does have those implications, but those implications are right. To see this, examine the mere events of the chair flying towards Phil and the chair hitting Phil. Consider each event by itself and the closest possible world where it doesn't occur. What happens in those worlds? What happens, I think, is that Knight tries to throw the chair at Phil and fails somehow. Maybe his aim is off and the chair misses Phil. Maybe his grip is bad and the chair tumbles off harmlessly. Whatever the case, the revised account's implications seem to depend on what Knight would do in response to things going wrong in such ways.

There are two possibilities here. Possibility one: Knight would do nothing. Maybe merely trying to hit Phil with the chair assuages Knight's anger. If so, the revised account doesn't have the implications that the objection claims. The account entails that the flying and the hitting are overall harmful for Phil because Phil would have been better off had they not occurred. Possibility two: Knight would have freely dismembered Phil. If so, the revised account does entail that the flying and the hitting are overall beneficial and merely pro tanto harmful for Phil. But that's the right result. If Knight would have freely dismembered Phil had those events not occurred, then they're overall beneficial because they *prevent* the dismemberment. By contrast, Knight's free act of chair throwing isn't like that. It doesn't prevent the dismemberment. That's partly why it's overall harmful. We can appeal to this difference to

³⁰ Thanks again to Eric Vogelstein for this objection.

justify the different assessments.³¹ The revised account gets the right results here.

You might think that this case can be modified to forestall this response, though. Suppose that Knight wouldn't miss because he's an expert chair thrower. So, if the flying or hitting hadn't occurred, it would have to have been because Knight didn't throw the chair and simply dismembered Phil instead. If so, the flying and the hitting don't prevent the dismemberment. They just correlate with its non-occurrence. So the objection resurfaces. Intuitively, the flying and the hitting are overall harmful for Phil in this variant of this case. But CCA counts them as overall beneficial and merely pro tanto harmful for Phil.

I agree that in this variant of the case the flying and the hitting are overall harmful for Phil. But the revised account can yield this result if it's combined with a plausible view of the distinction between free acts and mere events. We should say that the flying and the hitting don't count as mere events and that they're actually part of Knight's free act. This is because, given Knight's expertise at chair throwing, he brings these events about by exerting a degree of control that suffices for them to count as part of his free act. To illustrate: for the revised account's purposes, these events are sufficiently like my arm raising after I've willed it to raise under normal conditions. Under such conditions, my arm raising is part of my free act. Given Knight's chair-throwing expertise, the same goes for the relevant events in this case. So the events are covered by CCA+ and the revised account yields the right result about them: they're overall harmful for Phil.

To put the reply in general terms: the revised account can adequately deal with cases like this if it's combined with a view of free acts that uses the concept of control to distinguish free acts from mere events. No doubt, formulating such a view requires working out some important details. But I won't try to do that here. I've said enough to show that the objection at issue isn't obviously fatal to the revised account.

7. Conclusion

The omission and preemption objections to CCA don't raise obviously fatal problems for it. Certain preemption cases show that CCA can't adequately account for the harmfulness of free acts, given an important assumption about free will. Even if we don't want to reject this assumption, there's a promising solution to this problem. We should supplement CCA with

³¹ Note that appealing to this difference to justify the different assessments is consistent with my earlier argument for M. In Golf Clubs, there's no relevant difference between the Joker's actions and Batman's changing his mind. Both have the same effect on Robin. By contrast, there's a relevant difference between Knight's act and the mere events that follow: they affect Phil in different ways.

another principle that treats free acts differently from mere events. I've proposed such a principle and defended my revised account from some important objections.³²

References

- Boonin, D. (2014). *The Non-Identity Problem and the Ethics of Future People*, New York: Oxford University Press.
- Bradley, B. (2012). "Doing Away With Harm," *Philosophy and Phenomenological Research*, 85: 390–412.
- (2009). *Well-Being and Death*, New York: Oxford University Press.
- Brink, D. (2014). "Mill's Moral and Political Philosophy," *The Stanford Encyclopedia of Philosophy*, E. N. Zalta (ed.), <http://plato.stanford.edu/entries/mill-moral-political>.
- Feit, N. (2015). "Plural Harm," *Philosophy and Phenomenological Research*, 90: 361–88.
- Gardner, M. (forthcoming). "Beneficence and Procreation," *Philosophical Studies*.
- (2015). "A Harm-Based Solution to the Non-Identity Problem," *Ergo*, 17: 427–44.
- Hanser, M. (2008). "The Metaphysics of Harm," *Philosophy and Phenomenological Research*, 77: 421–50.
- Harman, E. (2009). "Harming as Causing Harm," in M. Roberts and D. Wasserman (eds.), *Harming Future Persons: Ethics, Genetics, and the Non-Identity Problem*, Dordrecht: Springer, pp. 137–54.
- (2004). "Can We Harm and Benefit in Creating?," *Philosophical Perspectives*, 18: 89–113.
- Kamm, F. (2014). "Responses," *Journal of Moral Philosophy*, 11: 476–517.
- Klocksien, J. (2012). "A Defense of the Counterfactual Comparative Account of Harm," *American Philosophical Quarterly*, 49: 285–300.
- Lewis, D. (1979). "Counterfactual Dependence and Time's Arrow," *Noûs*, 13: 455–76.
- (1973). *Counterfactuals*, Oxford: Blackwell Publishers.
- Mill, J.S. (1859). *On Liberty*, reprinted in M. Lerner (ed.), *Essential Works of John Stuart Mill*, New York: Bantam Books.

³² Thanks to audiences at the 31st annual Social Philosophy Conference, the 2nd annual Philosopher's Cocoon Philosophy Conference, the 2014 Pittsburgh Area Philosophy Colloquium, the 2014 Western Canadian Philosophical Association meeting, and the 2014 Central States Philosophical Association meeting. Thanks also to Ben Bradley for helpful discussion and to Alex Hyun and Lindsay Zi Lin, my commentators at the Philosopher's Cocoon Conference and the Central States Philosophical Association. And special thanks to an anonymous referee for this journal and to Eric Vogelstein, my commentator at the Pittsburgh Area Philosophy Colloquium. Both provided extensive comments that significantly improved the paper.

- Norcross, A. (2005). "Harming in Context," *Philosophical Studies*, 123: 149–73.
- Purves, D. (forthcoming). "Accounting for the Harm of Death," *Pacific Philosophical Quarterly*, doi: 10.1111/papq.12031.
- Rabenberg, M. (2015). "Harm," *Journal of Ethics and Social Philosophy*, www.jesp.org, vol 8, no. 3.
- Shiffrin, S. (2012). "Harm and Its Moral Significance," *Legal Theory*, 18: 357–98.
- (1999). "Wrongful Life, Procreative Responsibility, and the Significance of Harm," *Legal Theory*, 5: 117–48.
- Stalnaker, R. (1968). "A Theory of Conditionals," in N. Rescher (ed.), *Studies in Logical Theory*, Oxford: Blackwell.
- Tadros, V. (2014). "What Might Have Been," in J. Oberdiek, *Philosophical Foundations of the Law of Torts*, Oxford: Oxford University Press.
- Turner, P. N. (2014). "'Harm' and Mill's Harm Principle," *Ethics*, 124: 299–326.