Doing Away With Harm[1]

Ben Bradley

Syracuse University

Injunctions against harming others can be found everywhere. Here are a few prominent examples:

> *The Hippocratic Oath:* I will use those dietary regimens which will benefit my patients according to my greatest ability and judgement, and I will do no harm or injustice to them.[2]
>
> *J.S. 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).[3]
>
> *The Precautionary Principle:* Where an activity raises threats of harm to the environment or human health, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically.[4]

[1] Thanks to audiences at Lawrence University, the University of Wisconsin-Madison, SUNY-Buffalo, SUNY-Binghamton, Princeton University, Boston University, Syracuse University, Western Washington University, the 2009 Rocky Mountain Ethics Congress, and SPDMBABWS for helpful discussion. Here are just a few people who made helpful comments: David Killoren, Carolina Sartorio, Liz Harman, Melinda Roberts, Christopher Knapp, Daniel Star, Laurie Paul, Shieva Kleinschmidt, Christian Lee, Alastair Norcross, Eric Chwang, Chris Heathwood, Judith Thomson, Matt Skene, Kenneth Ehrenberg, Ken Shockley, Steve Campbell, Dave Horacek, Kris McDaniel, Melissa Frankel, Kara Richardson, Kevan Edwards. Special thanks to Justin Klocksiem and Brandon Davis-Shannon (my commentators at Boulder and Binghamton) and an anonymous referee for this journal.

[2] http://www.nlm.nih.gov/hmd/greek/greek_oath.html

[3] Holtug (2002) argues that we should reject the harm principle, in part because of difficulties with finding an account of harm that can be plausibly plugged into it. I believe Holtug is correct, and this paper might bolster his conclusion.

[4] http://www.gdrc.org/u-gov/precaution-3.html

Statements such as these suggest that there is something especially important about harm, such that we have strong, perhaps overriding reasons both to avoid harming people and to prevent harm from coming to people. Much of contemporary deontology is concerned with attempts to distinguish between doing and allowing or intending and foreseeing harm, and with justifying non-consequentialist constraints against harming others.[5] Despite the importance harm is supposed to have, almost nobody bothers to say what it is. This would not be a problem if harm were a primitive, undefinable notion, and if there were no significant disagreements about what counts as a harm. But harm is *not* plausibly a primitive, undefinable notion. And there *are* significant disagreements about what counts as a harm. So it is incumbent on philosophers to say what harm is.

Unfortunately, when we look at attempts to explain the nature of harm, we find a mess. The most widely discussed account, the comparative account, faces counterexamples that seem fatal. But no alternative account has gained any currency. My diagnosis is that the notion of harm is a Frankensteinian jumble. Thus it is unsuitable for use in serious moral theorizing. It should be replaced by other more well-behaved concepts, such as the axiological concepts of intrinsic and extrinsic badness. This conclusion will no doubt strike some as pessimistic. Given how little attention has been devoted to the concept of harm, isn't it possible that a plausible account will be forthcoming? In what follows I will discuss reasons to think that the answer is no.

Because so little work has been done in this area, there is groundclearing to do at the start, in the form of drawing key distinctions and setting out criteria for success.

1. Distinctions

[5] For a prominent recent example, see Kamm 2007.

The ontology of harm. In our ordinary talk about harm, we say that different sorts of things are harmful. Activities, like smoking, can harm. People can harm other people. Physical objects, like guns, can harm people. Perhaps other kinds of things (words?) can be harmful. For now, I will suppose that it is *events* that are harmful. When a person or a gun harms someone, it is in virtue of the person's or gun's involvement in some event that is harmful. (We will see below that there may be reason to revisit this decision, but it will not greatly affect any of the arguments to come.)

Intrinsic and extrinsic harm. Sometimes we say that smoking is harmful- attributing the property of harmfulness to smoking itself. Other times we say that smoking *causes* harm – implying that the harm is not the smoking itself, but something that happens later and is caused by the smoking, like the lung cancer. It makes sense that there are these two ways of talking about harm. In the theory of value we make a distinction between intrinsic value and extrinsic value. Something is intrinsically valuable just in case it is valuable in itself, or in virtue of its intrinsic properties. Something is extrinsically valuable if it brings about something else that is valuable, or prevents something bad. This distinction carries over straightforwardly to harm. Some events are *intrinsically* harmful: their mere occurrence constitutes harm to the person to whom it is occurring. Others are *extrinsically* harmful: they are harmful because of their effects. When we say that smoking is harmful, we are attributing extrinsic harmfulness to smoking. It is harmful in virtue of what it brings about, not because of what it is in itself.

When we say that pain is harmful we are saying that it is intrinsically harmful: just being in pain is harmful to the person experiencing it.[6]

In my view, to say that something is intrinsically harmful is just to say that it is intrinsically *bad* for the person undergoing it; it is to make a claim about well-being. What is intrinsically bad for someone is a matter of great controversy in the theory of well-being; some say pain, but others say desire frustration, or failure, or vice, or ignorance. Disputes about what makes for well-being are important, but not relevant to current disputes about the nature of harm. Thus I will focus on extrinsic harm. Others, however, might deny that intrinsic harm is tied to intrinsic badness in this way (see the end of Section 4 for discussion of this point). The arguments to come will not depend on identifying intrinsic harm with intrinsic badness.

Overall harm and pro tanto harm. There is another important distinction between types of harm. Sometimes when we attribute harmfulness to an event, we mean that it has some harmful feature, but we mean to leave open the possibility that it has other features that are beneficial and that outweigh the harmful features. In these instances we are attributing what we might call *pro tanto* harmfulness to the event. Other times we mean that, taking into account all the events harmful and beneficial features, the event is

[6] The distinction between intrinsic and extrinsic harms is similar to one drawn by Feinberg, who distinguishes between harmful acts and harmed conditions (Feinberg 1984: 31). Matthew Hanser seems to be following Feinberg when he distinguishes between *suffering* harm and *being* harmed: "I shall assume, provisionally, that being harmed is roughly equivalent to being caused to suffer it… Suffering harm is thus the more fundamental notion. And I take it that to suffer harm is simply to be its subject: connotations of pain and anguish should be ignored" (Hanser 2008: 421). But Hanser's way of defining what it is to suffer harm, "simply to be its subject," does not distinguish intrinsic from extrinsic harm, since one can be the subject of either sort.

on balance harmful. In these instances we are attributing all-things-considered or overall harmfulness. To illustrate, consider the following two headlines:[7]

Headline 1: **New study shows surgery is harmful!**

Headline 2: **Scientists develop harmless surgical technique!**

Headline 1 makes sense only if "harmful" is interpreted in an overall sense. We might imagine the ensuing story to describe a study showing that people who have surgery live shorter lives than patients with similar conditions who opt not to have surgery. If the story merely stated that patients who have surgery have their skin cut by scalpels, or experience pain afterwards, readers would feel justifiably misled by the headline and demand their money back. Headline 2 makes sense only if "harmless" is interpreted in a *pro tanto* sense.[8] We might imagine the story describing a new procedure where the surgery can be performed painlessly and without cutting into the patient's skin. If the story merely asserted that some patients are not, on the whole, worse off for having undergone this new procedure, readers would again feel cheated.

Overall harm and *pro tanto* harm are closely related. In fact they seem to be interdefinable. Whether something is all-things-considered harmful is a function of the ways in which it is *pro tanto* harmful or beneficial. We could define all-things-considered harm in terms of *pro tanto* harm in the following way: an event is overall

[7] Kagan gives a similar pair of examples designed to show that ordinary usage of 'harm' is split between a "global" use and a "local" use (Kagan 1998: 87). But he treats this as a conflict to be resolved, rather than as an indication of distinct concepts.

[8] Sometimes the distinction is drawn between all-things-considered harm and harm in a respect (Hanser 2008: 424, Kagan 1998: 86-7). I deviate from this use because I think the notion of "respects" of harm is contentious, since it seems to entail that *well-being* has different respects. Monists about well-being will deny this. An analysis of harm should not commit us to pluralism about well-being.

harmful to someone iff its *pro tanto* harms to that person outweigh its *pro tanto* benefits to that person. Alternatively, we could define *pro tanto* harm in terms of all-things-considered harm in the following way: an event is *pro tanto* harmful to someone iff it has an all-things-considered-harm-making feature.[9] I will focus on overall harm.

2. Desiderata

Now on to some desiderata for an analysis of extrinsic, overall harm. Some of these are desiderata for any sort of analysis; others are specifically important when analyzing harm.

a. *Extensional adequacy*. First and most obviously, the analysis must fit the data. If an analysis entails that I do not ordinarily harm someone by killing him, the analysis is false. If no analysis gets all the data right, we should favor the one that does better by the data, all else equal. To be extensionally adequate, an analysis should allow for harms of different degrees.

b. *Axiological neutrality*. Because we are interested in an analysis of *extrinsic* harm, not intrinsic harm, the analysis should not presuppose any substantive theory of well-being. For example, an analysis that entails that we can harm someone only by frustrating her desires should be rejected on the grounds that it presupposes a desire-based theory of well-being. Proponents of different axiologies should be able to agree--at some suitable level of abstraction--about what it takes for someone to be harmed, even if

[9] This might be unacceptable to certain extreme pluralists who deny that there is an all-encompassing sense of 'well-being'. They will claim that one can be harmed in this or that respect, but not all-things-considered. Even such pluralists need to employ a conception of all-things-considered harm, since it is possible to suffer both *pro tanto* harms and *pro tanto* benefits along a single dimension of well-being.

they might disagree about whether pain, or frustration, or something else, is required for harm.

c. *Ontological neutrality*. 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.

An acceptable analysis should also allow for different sorts of beings to be the subjects of harm. Examples in the literature tend to involve harms only to healthy adult human beings, but an analysis that entails that dogs or babies cannot be harmed must be false.

d. *Amorality*. The analysis should avoid moralistic fallacies. It should not presuppose that harming is morally wrong, or involves vicious intent.[10] The claim that harming actions are wrong is a substantive ethical claim; it isn't analytic. This condition follows from ontological neutrality (since there are non-actions that are harmful) but is worth independent emphasis. Suppose two individuals perform harming actions that have equal impact on their victims' well-being, but one does so with good intentions. This should plainly have no impact on whether the harms inflicted are of equal size; it affects only the blameworthiness of the agents or the wrongness of their actions. We wouldn't and shouldn't say that the agent whose intentions were good did not do harm, or

[10] Feinberg distinguishes a sense of 'harm' according to which x harm y iff x wrongs y (1984: 34-5). Like Feinberg, I am not concerned with this alleged sense of 'harm.'

did less harm, in virtue of his intentions. To ensure that facts about intentions or moral wrongness do not color our judgments, causing us to fall prey to moralistic fallacies, it is ideal to test analyses of harm by appeal to examples in which the harming event does not involve any intentionality at all.

e. *Unity*. The analysis should not merely be a list of some things that can happen to someone, nor should it have *ad hoc* features designed solely to account for particular cases. It should explain what *all* harms have in common by locating a common core to harm. Perhaps more controversially, it should also allow for a unified treatment of harm and *benefit*.

f. *Prudential importance*. The analysis should entail that harm is something worth caring about in prudential deliberation. Harm is the sort of thing we should try to avoid; if we have an analysis of harm such that one might reasonably be indifferent concerning whether an event of the sort described in the *definiens* takes place or not, we should reject the analysis. Or, more cautiously: we should either reject the analysis or give up on the idea that harm is an important concept in prudential deliberation.

g. *Normative importance*. Finally, the analysis should entail that harm is the sort of thing that it makes sense for there to be deontological restrictions about. If an analysis of harm, when plugged into Mill's harm principle or one of Frances Kamm's deontological principles, makes the principle absurd on its face, then it is not what we are looking for. The normative and prudential importance conditions might be rejected by someone who finds the notion of harm useless (Norcross 2005: 171-2).

This list is not meant to be exhaustive. Nor do I claim that any of these is an absolute requirement for an acceptable theory of harm; they are desirable features, but

perhaps it is impossible to meet them all. In the remainder of this paper I will introduce a simple comparative account of harm; I will spell out the problems this view faces; I will show how alternative views attempt to avoid those problems; I will show that those alternative views face their own, equally serious problems; I will sketch some attempts to save the comparative account from the apparently fatal problems; I will conclude that the comparative account cannot adequately capture our ordinary notion of harm, and that the ordinary notion should be jettisoned for purposes of philosophical theorizing.

3. The comparative account

The comparative account of harm is most often stated as a counterfactual account.[11] According to the counterfactual comparative account, a harmful event is an event that makes things go worse for someone, on the whole, than they would have gone if the event had not happened. The worse an event makes things go for someone, the more harmful it is. Conversely, a beneficial event is one that makes things go better for someone.[12]

The counterfactual comparison account fares well by most of the desiderata. It leaves open what sorts of beings may be harmed; any being that has a welfare can be harmed, not just a person. It also leaves open what sorts of events can be harmful; not just actions are harmful. It leaves open what welfare amounts to; it does not presuppose that pain, or desire frustration, or failure is intrinsically bad for people. Since it is

[11] A temporal comparative account, according to which an event harms someone if they are worse off afterwards than before, is sometimes presented as an alternative. For criticism of this view, see Holtug 2002: 368, Norcross 2005: 149-50, and Mora 2008: 38-40.
[12] See Feinberg 1984: 33-4, Parfit 1984: 69, Kagan 1998: 84, Norcross 2005: 150; Bradley 2009: 65.

ontologically and axiologically neutral, it is also amoral; it entails nothing about whether harming is wrong. It has explanatory power; it is not merely a list of harmful events, but provides a unified account of harms and benefits. It is based on an intuitively plausible idea: that harms *make a difference*, in a negative way, to the person harmed. And it seems that harm is rendered important, because we do indeed care to prevent or avoid events that make things go worse for us. It seems to get the right results in a great many cases. For example, if pain is intrinsically bad, and I now walk over to someone and kick him in the shins, the account correctly entails that my kicking is harmful to him in virtue of making him worse off. It also entails that death is harmful to anyone who has a good life to look forward to. So there are a lot of reasons to like this account of harm.

However, the comparative account appears to go wrong with respect to certain kinds of cases. I will discuss three.

First, we have a problem involving *preemption*. Suppose Batman drops dead of a heart attack. A millisecond after his death, his body is hit by a flaming cannonball. The cannonball would have killed Batman if he had still been alive. So the counterfactual account 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. So there seem to be cases where the account fails to count a harmful event as harmful.

Second, we have a problem involving *omission*. These are cases where the comparative account counts as harmful events that, intuitively, do not seem harmful, such as cases in which someone fails to receive some benefit. Failing to benefit someone

moves that person down on the well-being scale, and therefore counts as a harm. Suppose Batman purchases a set of golf clubs with the 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. The comparative account 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. So there are cases where non-harmful events are counted as harmful by the comparative account.

Third, we have the *non-identity problem* (Parfit 1984: Ch. 16). Suppose Mary is contemplating pregnancy. If she becomes pregnant now, she will conceive a child, Jane, who will have a painful disease. If she waits a few months to conceive, she will conceive a different child, John, who will not have that disease. In that case, Jane would never come into existence at all. Mary chooses to conceive Jane. Jane lives a good life on the whole, despite the pain she endures from her disease; but due to all that pain, her life is much worse than the relatively pain-free life John would have had if she had waited. It seems Mary acts wrongly, and the wrongness of her act is explained by the harm her act inflicts on Jane. But the comparative account seems incompatible with the claim that Jane is harmed by being brought into existence, because coming into existence brings more benefit than harm to Jane. All things considered, Jane is better off as a result of Mary's bad decision.

The problems about preemption and omission strike me as very serious. The non-identity problem strikes me as less serious; I don't think Jane is harmed. All three

problems have been taken to be fatal to the comparative account. I will examine attempts to defend the comparative account, but first will look at some alternatives.

4. Non-comparative Accounts

Elizabeth Harman is concerned about the implications of the comparative account in non-identity cases. Harman offers the following as a sufficient condition for an act to be harmful: "one harms someone if one causes him pain, mental or physical discomfort, disease, deformity, disability, or death" (Harman 2009: 139). Let us generalize this account to be an account of harmful events more generally:

> (1) An event harms someone if it causes him pain, mental or physical discomfort, disease, deformity, disability, or death. (Harman 2009: 139)

(1) is what we may call a non-comparative account, since according to it, whether an event is harmful to someone does not depend on a comparison with how things would have gone for that person otherwise. (1) has some nice features. If (1) is true, Mary harms Jane by causing her to have the disease, even though Jane would not have existed if Mary had not caused her to have that disease. So we would have a solution to the non-identity problem. (1) also seems to avoid the other problems facing the comparative account. It avoids the preemption problem; since Batman's heart attack causes him to die, it doesn't matter that the cannonball would have caused him to die; the heart attack still harms him, according to (1). It also avoids the omission problem, since omitting a

benefit need not bring about any of the conditions on Harman's list (and if an omission were to bring about one of those conditions, it would seem to be a genuinely harmful omission).

But (1) is hardly an analysis of harm. It cries out for unification- what do the items on the list have in common? To see further why this is a problem, note that (1) provides no way to determine degrees of harm. It tells us that events causing disease and injury are harms, but does not tell us how harmful those events are. Fortunately, Harman does not take her list to be end of the story, but offers the following more general statement of sufficient conditions for harm: "an action harms someone if it causes the person to be in a bad state. Bad states are understood as states that are in themselves bad, not bad because they are worse than the state the person would otherwise have been in" (2009: 139). Summing this up, we have Harman's sufficient condition (again generalized from actions to events generally):

H. An event harms someone if it causes the person to be in an intrinsically bad state.

We may wonder whether H is a fair generalization of (1). After all, while pain and discomfort are plausibly intrinsically bad, the other items on Harman's list are not. (Some of these might turn out to be intrinsically bad on some perfectionist accounts of well-being; but those accounts are not plausible.) Still, H is a view worth considering in its own right.

H is not a competitor to the comparative account. First, H is not a statement of necessary and sufficient conditions; it is merely a statement of sufficient conditions. So

we cannot assess its extensional adequacy on a par with that of the comparative account. We can test it for false positives, but not false negatives.

Second, H appears to be an account of *pro tanto* harm rather than all-things-considered harm.[13] But if it is an account of *pro tanto* harm, then it is compatible with the comparative account, since they are about different things. Defenders of the comparative account will no doubt agree that events that cause intrinsically bad things to happen to people are at least *pro tanto* harmful.

Here is an extended passage from Seana Shiffrin in which she sketches her account of harm:

> On my view, harm involves conditions that generate a significant chasm or conflict between one's will and one's experience, one's life more broadly understood, or one's circumstances. Although harms differ from one another in various ways, all have in common that they render agents or a significant or close aspect of their lived experience like that of an endurer as opposed to that of an active agent, genuinely engaged with her circumstances, who selects, or endorses and identifies with, the main components of her life. Typically, harm involves the imposition of a state or condition that directly or indirectly obstructs, prevents, frustrates, or undoes an agent's cognizant interaction with her circumstances and her efforts to fashion a life within them that is distinctively and authentically hers—as more than merely that which must be watched, marked, endured or undergone. To be harmed primarily involves the imposition of conditions from which the person undergoing them is

[13] Harman appears to contrast her account with an "all-things-considered" account (2009: 141).

> reasonably alienated or which are strongly at odds with the conditions she would rationally will" (Shiffrin 1999: 123-4).

Shiffrin goes on to list some of the things that count as harms on her account: pain, disability, illness, injury, and death (124). The reason these things are harmful is that they are states that a person rationally wills not to be in. It seems fair to sum up Shiffrin's view in the following way:

> S. An event harms someone iff it causes her to be in a state that she rationally wills not to be in.

It seems like a mistake to analyze harm in terms of what one rationally wills, for two reasons. (1) In general, it seems possible to rationally will that one be harmed – for altruistic reasons, say – but Shiffrin's view rules this out. (2) If only creatures that rationally will things, or try to fashion lives that are distinctively and authentically theirs, can be harmed, then many non-human animals, human babies, and certain brain-damaged human beings cannot be harmed at all. This seems wrong, and violates ontological neutrality.[14]

We can solve these problems by removing the talk of rational willing, and amending Shiffrin's account as follows:

> S'. An event harms someone iff it causes something intrinsically bad to happen to him.

[14] For more criticism of Shiffrin's view see Mora 2008: 105.

I have no reason to think Shiffrin would accept S'. S' is just Harman's view made into a biconditional; according to this view, causing something intrinsically bad to happen to someone is not merely sufficient, but also necessary, for harm.

This view faces a problem that any non-comparative account will face: how to account for preventive harms, and in particular the harm of death. Death—i.e. the event of one's death—causes one not to exist, or to become a corpse. It is not *intrinsically* bad or harmful, nor does it cause anything intrinsically bad for its victim.

One might worry that this critique of non-comparative accounts is unfair. Perhaps Harman or Shiffrin could say that they are really interested not in giving accounts of what it is for an event to bring about harm, but of what it is to be in an intrinsically *harmed state,* where this is *not* to be understood as being in an intrinsically bad state. (If this is Harman's view, it would be more understandable that she takes (H) to be a generalization of (1) despite the fact that death is not intrinsically bad.) Perhaps to be dead is to be the subject of this sort of harm. But even if we can make sense of being in a harmed state while not existing, we still need to know what makes it the case that being dead is being in a harmed state. On a non-comparative account, we cannot appeal to the lost goods of life to explain this, and thus we cannot account for the harm of death, for if death is harmful, it must be in virtue of what it prevents its victim from having. Prevention is most naturally construed comparatively. Thus all non-comparative accounts lack the resources to consider death – and, by extension, events that cause death, such as killings – to be harmful.

5. Non-consequentialist Harm?

Some have distinguished between "consequentialist" and "non-consequentialist" accounts of harm. Here is James Woodward: "On a non-consequentialist approach we think of a person as harmed… whenever an action is performed that violates some right possessed by or obligation owed to that person," while on a consequentialist approach we think of harm as having to do with well-being (1986: 818). Woodward describes some cases that appear to be instances of harm, but would not qualify as harms on the comparative account. Suppose Batman attempts to purchase an airline ticket, but the airline is religiously intolerant, allowing only Christians to fly. Batman is turned away by the airline worker due to his membership in the Church of the Subgenius. As it happens, the plane he was attempting to get on crashes, killing all aboard. The comparative account entails that the airline worker did not harm him, since he is better off as a result of being denied a ticket; but according to Woodward, Batman was still harmed by the airline worker's discriminatory act, because his rights were violated.[15]

There are different ways we might take Woodward's point. He might be arguing that consequentialist and non-consequentialist accounts of harm are competing accounts, and that we should adopt a non-consequentialist account of harm. But this seems highly implausible. The non-consequentialist account offered by Woodward cannot account for the harmfulness of any non-actions, such as explosions, earthquakes, and rises in global temperature. None of these events violate rights, but nevertheless they are very harmful to people. Woodward's account also cannot account for the harmfulness of punitive acts

[15] See Woodward 1986: 810-11. Some details have been changed. For more discussion of Woodward's view, see Mora 2008: 59-60.

performed against people who have forfeited the relevant rights. If someone takes away Robin's money, this is harmful to Robin, even if he stole the money and therefore has no right to it. If Woodward's account of harm is supposed to *compete* with the comparative account, it will lose.

Perhaps the point, instead, is that there are different concepts answering to the name 'harm': a consequentialist one and a non-consequentialist one.[16] I don't want to rule out this possibility. But it strikes me as misleading to use the term 'harm' to refer to this non-consequentialist notion. From the standpoint of moral theory, we could dispense with talk about non-consequentialist harm and replace it with talk about rights violations, thereby losing nothing of substance, while avoiding the confusion that would likely result from employing two concepts that go by the same name.

6. Hanser's Event-Based Account

Matthew Hanser sees the problems with extant comparative and non-comparative accounts as sufficiently grave to warrant an entirely new approach, which he calls the "event-based" account. Here are some crucial definitions:

> The account begins with the notion of a basic good. 'Goods' are not states or conditions that it is good to be in. Rather, they are things that it is good to have. And 'basic' goods are, roughly speaking, those the possession of which makes possible the achievement of a wide variety of the potential components of a reasonably happy

[16] Perhaps these concepts would correspond to the second and third senses of 'harm' distinguished by Feinberg (1984: 33-5), where the second appeals to interests and the third appeals to rights.

> life… I shall provisionally take the basic goods to include certain fairly general physical and mental powers or abilities. (2008: 440-1)

Hanser develops his view using the notion of "levels" of harm and benefit, where a base-level harm is a loss of a basic good, a base-level benefit is an acquisition of a basic good, and higher-level harms and benefits are preventions of lower-level benefits and harms, respectively. Thus, Hanser's event-based account of harm is the following:

> HEBA: Someone suffers a harm if and only if he suffers a harm of some level with respect to some basic good. (2008: 442)

The motivation for Hanser's account is to give an account that avoids problems with preemption and accounts for the harm of death, which he takes to be the main problems facing comparative and non-comparative accounts. Hanser thinks HEBA avoids the preemption problem, because no matter whether the heart attack or the flaming cannonball kills Batman, he suffers a loss of basic goods, and is therefore harmed. But this alleged advantage of Hanser's view is merely the result of the fact that HEBA is answering a different question from the one answered by the comparative account. HEBA tells us under what conditions a person suffers harm. It does not tell us under what conditions an event is harmful to someone. A natural move for Hanser would be to say that an event is harmful iff it causes someone to suffer harm (where suffering harm is understood in accord with HEBA). Assuming a non-contrastive notion of causation, this would result in another version of a non-comparative account, rather than a third option.

It is the non-comparative account that does the work in avoiding the preemption problem, not HEBA. But the combination of HEBA with a non-comparative causation-based account of harming events might do the two things Hanser wants: the non-comparative account handles the preemption problem, while HEBA allows the view to account for the harm of death, since death involves a loss of basic goods.

But HEBA fails for other reasons. According to HEBA, whether someone suffers harm depends on whether the person loses a power, specifically the power to achieve things that are intrinsically good; losing the power is merely extrinsically bad for the person, since it results in the person missing out on the possibility of getting those intrinsic goods. Harm is analyzed by appeal to loss of extrinsic goods, not by appeal to loss of intrinsic goods. This suggests that something has gone wrong, or at least that the full story has not yet been told. We care about extrinsic goods only insofar as we care about the intrinsic goods they lead to, or the intrinsic evils they prevent. If we want to know *how* harmful it is for someone to lose a power, we have to look at the intrinsic values of the goods the power would have brought about and the evils it would have prevented. Without looking at those intrinsic values, how can we explain why it would be a worse harm to be paralyzed for ten years than to lose one's sense of smell for ten minutes? Hanser suggests that some basic goods are "more important" than others (2008: 444), but how is this to be understood without appeal to the intrinsic goods to which they lead? Surely it is not a brute fact that some extrinsic goods are more important than others. But if those intrinsic values are relevant to the harmfulness of losing a power, then we are back to a comparative account of the sort Hanser was trying to avoid.

Furthermore, as an account of what it is to *suffer* (rather than cause) harm, Hanser's account seems to fail. Hanser analyzes harm in solely in terms of losses of basic goods. But there are some cases of suffering harm that do not involve any losses of basic goods: for example, pains. Suffering great pain need not count as suffering harm, according to HEBA, since it need not involve the loss of any capacities.[17]

Hanser raises another problem for the comparative account, and says that his account avoids it. He suggests that according to the comparative account, harm is a relational notion: a person may be in a harmed state relative to one event, and a distinct harmed state relative to a different event.

> When someone would have been better off had a certain event not occurred, then, let us say (i) that he is in a harmed state relative to that event's occurrence, and (ii) that the event comes to him as a harm. It is important to see that according to this account, the notion of harm is relational. If someone would have been better off had an event E1 not occurred and also had a distinct event E2 not occurred, then he is in two distinct harmed states: that of being worse off than he would have been had E1 not occurred and that of being worse off than he would have been had E2 not occurred. (2008: 423)

Hanser's way of describing the comparative account, as involving a person being in a harmed state relative to an event, is convoluted, in something like the way it would be convoluted to say that I am in a married state relative to my wife. The comparative

[17] Hanser briefly mentions the possibility of basic bads, but pain does not seem to fit the bill (it is not the sort of thing one can "acquire"; it is not analogous to, say, vision or the ability to walk; it is intrinsically bad, not a power to bring about intrinsic bads).

account states that a harming event makes a person worse off than she would have been; as I understand it, the account does not explain the harmfulness of an event by appeal to harmed states at all. If all Hanser means by being "in a harmed state relative to that event's occurrence" is being such that it is *that* event, rather than some other event, that made her worse off, this is unobjectionable. But the talk of "harmed states" naturally leads one to focus on intrinsic states of the person. The comparative account denies that whether an event harms someone is just a matter of the intrinsic properties of that person. It is an account of extrinsic harm, not intrinsic harm.

This is important, because Hanser goes on to argue against the comparative account on the grounds that it is implausible to suppose that someone could be in a distinct harmed state relative to each harming event. He writes:

> Suppose that A shoots B, causing him to become paralyzed from the waist down. A whole series of causally linked events occur here, among which are A's pulling of the trigger, the gun's going off, the bullets entering B's body, and B's becoming paralyzed. According to the counterfactual comparison account, B is in a distinct harmed state relative to each of these events… But I think this is clearly the wrong way to describe the situation. What we have here are not four separate harms, but a single harm—B's becoming paralyzed—with multiple causal antecedents. The counterfactual comparison account collapses the distinction between events that cause people to suffer harms and the harms that the people are thereby caused to suffer. (2008: 433)

In Hanser's example, there are indeed four extrinsically harmful events. There is perhaps only one intrinsic harm: B's paralysis. (I doubt this is actually an *intrinsic* harm, since on my view intrinsic harms are intrinsic bads, and it is probably not intrinsically bad to be paralyzed. But let this pass for now; being paralyzed is, on Hanser's account, a base-level harm). What, then, is the objection supposed to be? Is the objection that distinct elements of a causal chain each turn out to be harming events? This might be problematic if we thought that the total harm produced by the four-event sequence were equal to the sum of the harmfulness of the four elements taken individually; but nobody thinks the harmfulness of extrinsically harmful events should be added in this way. It seems like an objection because Hanser takes the comparative account to be an account of what it is to be in a harmed state. And it doesn't seem like B *suffers* four distinct harms, or is in four distinct harmed states. But that is because talking about harmed states, or harms suffered, leads us to think of intrinsic states of the victim, and B is not in four distinct intrinsic states corresponding to different harms. The counterfactual comparative account is not an account of intrinsic harm, so Hanser's objection leaves untouched the comparative account of harm discussed here.

7. Can the comparative account be rescued?

Since there seems to be no better alternative on offer, we should return to the comparative account to see how it might be defended against the other arguments. First, the non-identity problem: is this a problem at all? Recall the distinction between all-things-considered harm and *pro tanto* harm. In the example of Mary and Jane, we may

be able to say that Jane suffers a *pro tanto* harm by being brought into existence. Bringing her into existence causes her some pain, and causing pain is the sort of thing that tends to make an event all-things-considered harmful. So if we want to explain the wrongness of Mary's decision by appeal to harm to Jane, we might appeal to its *pro tanto* harmfulness.

But suppose the case is slightly changed. Suppose that instead of being born with a painful illness, Jane will be born with a diminished ability to enjoy life. Again, her life will be good on the whole, but much less good than the life John would have had if Mary had waited a few months to conceive. It still seems Mary ought to wait. But Jane is not harmed in any way by Mary's decision to go ahead and conceive. She is caused no pain, and is deprived of nothing good, by Mary's decision.

In my view, we should just deny that the wrongness of Mary's decision is explained by its harmfulness to Jane. Someone who owes her existence to a decision that left her with a painful disease may be benefited by that decision if her existence is overall worthwhile; she is not harmed, all things considered. She might suffer a *pro tanto* harm, but it is outweighed by the *pro tanto* benefits conferred by the decision. The *pro tanto* harms she suffers do not make the decision overall harmful to her, but they do affect the moral evaluation of the act, since there was another alternative available to the mother that would not have been *pro tanto* harmful in that way. The wrongness of the act does not follow from its harmfulness to the child, but from the existence of the better alternative.[18] (This is of course a deep and difficult issue; see Roberts and Wasserman (eds.) 2009 for recent discussion.)

[18] It is also possible to argue that although the child is not harmed, the child is nevertheless wronged; one can wrong without harming by, for example, violating rights.

Next consider the preemption and omission problems. Perhaps, just as we should revise our judgments about harmfulness in non-identity cases, we should revise our judgments in preemption cases. Suppose Batman knew he was going to get hit by the flaming cannonball. How much would he care if he dropped dead of a heart attack a millisecond before? Probably not at all; nor should he, it seems to me. And if he doesn't and shouldn't care about it, this gives us some reason to say it isn't harmful, at least if we want harm to be something we care about. So there is some plausibility to this strategy.

I think, however, that refusing to call any such preempting events harmful does great violence to our ordinary practice of harm attribution. Alastair Norcross describes a case that, briefly, goes as follows: Bobby Knight chokes a philosopher, injuring her windpipe; if he hadn't choked her, he would have torn her arms off, which would have been much worse for her. The comparative account entails that Knight did not harm her, but benefited her, by choking her (Norcross 2005: 165-66). This case seems to show that according to the comparative account of harm, one can make one's pain-causing actions overall beneficial rather than harmful merely by 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. As in the cannonball case, we might say: all things considered, shouldn't the victim prefer that the choking happen rather than not, and if so, doesn't that indicate that the choking is not all-things-considered harmful? Perhaps; but the judgment that serious harm has occurred does not evaporate. There is, of course, *pro tanto* harm in this case, in the form of the crushed windpipe and its attendant pain. The counterfactual account is compatible with this. But it seems that there is serious overall harm here too.

There are at least three ways that we might try to avoid the preemption problem within the framework of a comparative account. There is a contrastive account, a maximizing account, and a causal account. I consider each in turn.[19]

First, the contrastive account. I have supposed, in formulating the comparative account, that harm is a relation between an event and an individual who is harmed. But perhaps this picture is too simple. In the Knight case, we may want to say two things at the same time: Knight's choking the philosopher, *rather than merely saying "hi,"* is harmful to her. Knight's choking the philosopher, *rather than ripping off her arms*, is beneficial to her. Perhaps, then, harm should not be taken to be a relation between an event and an individual, but a relation between an event, a contrast event, and an individual. So the fundamental harm-facts would be complex states of affairs of the form *E1 rather than E2* harms S, where E2 is a distinct alternative event that could have happened instead of E1.[20] The very same event, such as Knight's choking the philosopher, can be involved in many such complexes involving different alternative events. Ordinary harm-attributions often fail to contain an explicit reference to a contrast-event; in such cases their meanings may be indeterminate, or may be fixed by features of conversational context. This is the *contrastive account* of harm.

Unfortunately, the contrastive account fails to make harm normatively important. For just about any act A1 we might think of, there will be alternative acts A2 and A3, such that the agent's doing A1 rather than A2 harms someone while the agent's doing A1

[19] In a paper forthcoming in this journal, Judith Thomson defends a different sort of comparative account (Thomson forthcoming). It is in some ways similar to the causal account, but in other ways different. I hope to be able to discuss Thomson's new account in future work.

[20] This account is inspired by the contrastive accounts of causation championed by, e.g., Christopher Hitchcock (1996) and Jonathan Schaffer (2005). Norcross's comparative account of harm is a contrastive account (2005: 167-8).

rather than A3 does not. *But an agent cannot do A1 rather than A2, without also doing A1 rather than A3.* The agent can merely choose to do A1. And on the contrastive account, A1 itself is not harmful (even if we can give a semantics for 'harm' according to which statements attributing harmfulness to an act, without explicitly mentioning a contrast event, turn out true). Thus, no deontological principle prohibiting harm will be remotely plausible if the contrastive account of harm is true.[21]

Another comparative view is given by Melinda Roberts. According to Roberts's view, whether an act is harmful to S depends not on what would have happened if the act had not occurred, but on whether the agent had an alternative that would have been better for S.[22] Thus, in the Knight case, since Knight had a better alternative available, such as saying "hi," his choking act counts as harmful even though he would have done something worse had he not done that. His violent temper does not diminish the harmfulness of his acts.

Roberts's view does not get around the problem of omitting benefits. If a surgeon saves someone's life, he does not harm the patient by failing to put $20 in the patient's pocket at the end of the procedure. But the alternative where he puts the $20 in is better for the patient than the alternative where he doesn't. So Roberts's view seems to entail that the surgeon harms the patient.

Roberts's view has an additional problem. Since it explains harm by appeal to alternative actions, it seems to apply only to actions performed by autonomous agents,

[21] Since Norcross is not a deontologist, he would not consider this to be an objection. The goal here, however, is to see whether there is a *useful* notion of harm.
[22] Roberts 1998: 32, Roberts 2009.

not to natural events.[23] It is not clear how the account might be extended to account for the harmfulness of earthquakes. So Roberts's view requires us to treat actions and non-actions differently, violating ontological neutrality.

Finally, we might try to avoid some of the problems of the counterfactual account by explaining harm in terms of causation rather than counterfactuals. Instead of comparing what happens with what would have happened, we could compare what an event *causes to happen* with what it *causes not to happen* (Bradley 2004; Conee 2006). It is natural to think that the heart attack causes Batman to die, even though he would have died anyway from a different cause; it causes him to be deprived of a valuable future, even though the cannonball would have done so if it hadn't. An event can cause something to happen (or not to happen), even if it would (or wouldn't) have happened anyway. Call everything that an event causes its *total causal consequence*; call everything that an event causes not to happen its *total prevention*. According to the comparative account I am suggesting, the causal comparative account, an event is harmful to someone if and only if its total causal consequence is worse for that person than its total prevention. The causal account entails that Batman's heart attack is bad for him. Its total causal consequence has zero value for him, while its total prevention has very high value for him, since it includes all the good things that the heart attack prevents him from getting.

Thus the causal account *seems* closer to being extensionally adequate than the counterfactual account. But it is not clear that the preemption problem is solved. For

[23] Roberts is aware of this problem and claims that her account is only supposed to cover culpable harm; it is not supposed to capture the "ordinary language sense" of the term (1998: 176n52). But it seems to me that if 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. We are using 'harm' in its ordinary language sense whether we talk about harm by falling boulders or harm by boulder-pushing acts.

nothing has yet been said about which are the events that are caused not to happen by an event. If the events that Batman's heart attack causes not to happen are just the ones that would have happened otherwise, then we are back to a counterfactual account. If other events, which ones? The heart attack seems to cause Batman not to have any more good experiences (even though he wouldn't have had them anyway). Does it also cause him not to become President of the U.S.? He wouldn't have become President even if he hadn't died from the heart attack, but this is irrelevant, since we are supposing that events may be caused not to happen without being such that they would otherwise have happened. Moreover, a causal account along these lines does not appear to help with the Bobby Knight case, where the actual harm and the preempted harm are to different parts of the body. Knight's choking the student causes him not to rip off her arms. Thus its total causal consequence seems to have higher value for the student than its total prevention, making the choking beneficial for the student.

8. Upshots

We face the following dilemma. When we focus on overall extrinsic harm, we are inevitably led to a comparative account, since no non-comparative account offers a way to account for preventive or deprivational harms such as the harm of death. Non-comparative accounts are plausible only as partial accounts of *pro tanto* harm. But comparative accounts are not fully satisfactory either. The counterfactual account has problems with preemption and omission. Roberts's maximizing account has problems with overcounting harms of omission, and is not ontologically neutral, since it requires us

to treat harmful acts and harmful non-acts differently. The causal account lacks the resources to prevent overcounting of preventions, and does not solve all preemption problems anyway. The contrastive account avoids at least some of those problems, but makes harm theoretically useless.

Why is it so difficult to give an acceptable account of harm? I suspect part of the problem might have to do with moralizing.[24] We don't say I harm all the people to whom I do not give my money, thereby not benefiting them, thereby leaving them worse off than they would have been if I had given them the money. We do say I harm my children if I fail to provide for them. Why? Perhaps because we think it is impermissible to fail to provide for my children, but permissible to fail to provide for others' children. We are more likely to call an act harmful if we think it is wrong. Perhaps this impermissibility has to do with rights, justice, or desert; my children deserve my assistance, they have a right to it, whereas strangers do not; I (usually) do not *wrong* a stranger by failing to give them money that would help them. If this diagnosis is accurate, our judgments about harm are muddied by moralizing. I do not claim this is a complete explanation for the problems involved in giving an account of harm, but it seems plausible to suppose that it plays a role.

Suppose my arguments have been persuasive. What are the prospects for deontological principles forbidding harm? Not good. A comparative account of harm seems required in order to account for the harm of death – presumably the most important harm for which deontological principles need to account! But comparative accounts allow for many more harms than a deontologist could possibly want to prohibit, such as

[24] Holtug argues that we get an account of harm that is more suitable (but not wholly suitable) for employment in Mill's harm principle if we moralize harm, so that harm always involves wrongdoing (2002: 377-80).

harms of omission. The contrastive account seems to be the best of the comparative accounts, but as we have seen, it is utterly useless for deontology. So there seems to be an important challenge for deontological ethics here.

It seems to me – though I have obviously not established this conclusively – that the best course of action is to avoid appealing to the notion of harm at all in our moral theorizing. Consequentialists don't need it; they can appeal to the intrinsic values of consequences of acts. Deontologists don't need it; they can appeal to the other notions they think are relevant to wrongness, such as the violation of rights or the inflicting of pain or injury. Nobody really needs to talk about harm, and doing so invites unnecessary confusion. Let harm go the way of phlogiston.

References

Bradley, B. (2004) "When Is Death Bad for the One Who Dies?" *Nous* 38: 1-28.

Bradley, B. (2009) *Well-Being and Death*. Oxford: Clarendon Press.

Conee, E. (2006) "Dispositions toward Counterfactuals in Ethics," in McDaniel et al. (eds), *The Good, the Right, Life and Death* (Ashgate), pp. 173-87.

Feinberg, J. (1984) *Harm to Others*. Oxford.

Hanser, M. (2008) "The Metaphysics of Harm." *Philosophy and Phenomenological Research* 77: 421-50.

Harman, E. (2009) "Harming as Causing Harm," in Roberts and Wasserman (eds.), *Harming Future Persons*, pp. 137-54.

Hitchcock, C. (1996) "Farewell to Binary Causation." *Canadian Journal of Philosophy* 26: 267-82.

Holtug, N. (2002) "The Harm Principle." *Ethical Theory and Moral Practice* 5: 357-89.

Kagan, S. (1998) *Normative Ethics*. Westview.

Kamm, F. (2007) *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford.

Mill, J.S. (1859) *On Liberty*. Reprinted in M. Lerner (ed.), *Essential Works of John Stuart Mill* (New York: Bantam Books, 1961).

Mora, E. (2008) "Understanding Harming Actions." M.A. Thesis, Monash University.

Norcross, A. (2005) "Harming in Context." *Philosophical Studies* 123: 149-73.

Parfit, D. (1984) *Reasons and Persons.* Oxford.

Roberts, M. (1998) *Child versus Childmaker: Future Persons and Present Duties in Ethics and the Law*. Lanham, MD: Rowman and Littlefield.

Roberts, M. (2009) "The Nonidentity Problem and the Two-Envelope Problem: When Is One Act Better for a Person than Another?" in Roberts and Wasserman (eds.), *Harming Future Persons*, pp. 201-228.

Roberts, M. and Wasserman, D. (eds.) (2009) *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*. Springer.

Schaffer, J. (2005) "Contrastive Causation." *The Philosophical Review* 114: 327-58.

Shiffrin, S. (1999) "Wrongful Life, Procreative Responsibility, and the Significance of Harm." *Legal Theory* 5: 117-148.

Thomson, J. (forthcoming) "More on the Metaphysics of Harm." *Philosophy and Phenomenological Research*.

Wingspread Statement. (1998) Available at http://www.gdrc.org/u-gov/precaution-3.html.

Woodward, J. (1986) "The Non-Identity Problem." *Ethics* 96: 804-31.