

The Version of Record of this manuscript has been published and is
freely available in the Australasian Journal of Philosophy, 2022, [http://www.tandfonline.com/
10.1080/00048402.2022.2048401](http://www.tandfonline.com/10.1080/00048402.2022.2048401).

A HYBRID ACCOUNT OF HARM

Charlotte Franziska Unruh

Abstract

When does a state of affairs constitute a harm to someone? Comparative accounts say that being worse off constitutes harm. The temporal version of the comparative account is seldom taken seriously, due to apparently fatal counterexamples. I defend the temporal version against these counterexamples and show that it is in fact more plausible than the prominent counterfactual version of the account. Non-comparative accounts say that being badly off constitutes harm. However, neither the temporal comparative account nor the non-comparative account can correctly classify all harms. I argue that we should combine them into a hybrid account of harm. The hybrid account is extensionally adequate and presents a unified view on the nature of harm.

Keywords

Harm, Hybrid Account, Well-being, Temporal comparative harm, Counterfactual comparative harm, Non-comparative harm

1. Introduction

Two questions are of central importance in the philosophy of harm. First, under which conditions can we say that someone suffers a *harm*, in the sense of being in a harmed state?¹ Second, under which conditions can we say that one person *harms* another, in the sense of bringing about the harmed state? To illustrate, imagine that Ann throws a stone at Bob, thereby breaking his nose. We can now ask whether, and if so, in virtue of which facts, Bob is

¹ In the following, I will use ‘suffering harm’ and ‘being in a harmed state’ synonymously, without meaning to suggest that harm is constituted by suffering or that harm necessarily comes with the felt experience of suffering.

in a harmed state.² We can also ask whether, and if so, in virtue of which facts, Ann has harmed Bob.

These questions are not always clearly distinguished in the literature.³ As a result, it is easy to overlook the fact that the notion of harm is distinct from the notion of harming and deserves separate attention. Perhaps the focus on harming is not surprising, since the direct objects of moral and legal prescriptions are behaviours [Feinberg 1984: 31]. However, it would be a mistake to think that we should pay less attention to the notion of harm. Indeed, it seems plausible that any account of harming needs to presuppose an account of harm (see Feinberg [1984: 31]). Intuitively, for a behaviour (such as Ann's throwing the stone) to count as harming, the behaviour needs to be related, in an appropriate way, to an outcome that counts as a harm (such as Bob's broken nose). An account of harm, then, can help us to identify and compare harms, specify when moral principles against harming apply, and ideally, explain what unites different instances of harm and what makes them morally significant.

In this paper, I develop a novel hybrid account of harm and defend it against the main accounts of harm in the literature, the comparative account and the non-comparative account:

(Comparative) Agent suffers a harm if and only if Agent is worse off.

(Non-Comparative) Agent suffers a harm if and only if Agent is badly off.⁴

(Hybrid) Agent suffers a harm if and only if Agent is (i) worse off, or (ii) badly off.⁵

² I assume that we should conceive of harms as states of affairs, rather than events (for development and defence of an event-based view on harm, see Hanser [2008]).

³ A noteworthy exception is Hanser [2019: 853]. Similar points about the relation between harm and harming have also been made by Feinberg [1984: 31], Hanser [1990: 57], and Shiffrin [2012: 360].

⁴ Accounts along these lines have been defended by Harman [2004, 2009], Shiffrin [1999], and Rivera-lópez [2009].

⁵ While some have suggested that an account of harm might have to combine comparative and non-comparative elements (McMahan [2013: 8, n. 3], Woollard [2012: 688], the possibilities of such an account have not been comprehensively explored. An exception is Meyer's [2016] disjunctive notion of harm, which I discuss in footnote 21.

There are different versions of the comparative account. According to the *temporal* account, Agent suffers a harm if and only if Agent is worse off than Agent was previously.⁶ The temporal account is seldom taken seriously in the literature, due to apparently fatal counterexamples. The much more prominent version of the comparative account is the *counterfactual* account, which says that Agent suffers a harm if and only if there is an event such that Agent is worse off than Agent would have been in the absence of this event.⁷

In this paper, I argue that taking the distinction between harm and harming seriously enables us to respond to apparent counterexamples to the temporal comparative account and provides a reason to reject the much more prominent counterfactual comparative account of harm. However, neither comparative nor non-comparative accounts can correctly classify all harms. I therefore propose that we should combine the temporal comparative account and the non-comparative account into a hybrid account. Temporal comparative harms consist in loss of well-being. Non-comparative harms consist in the presence of ill-being. Unlike its competitors, my hybrid account correctly classifies all harms and can be presented as a unified account of the nature of harm and benefit.

I borrow desiderata for accounts of harm from Bradley [2012: 394-95]. First, accounts of harm should be extensionally adequate. They should not identify states as harms that are clearly not instances of harm, but they should also not miss out states that clearly are instances of harm. Second, accounts of harm should be axiologically neutral. They should be compatible with different accounts of welfare. Third, accounts of harm should be unified, in the sense that they ‘should explain what all harms have in common’ [Bradley 2012: 395].⁸

⁶ A version of the temporal comparative account of harm is defended by Velleman [2008: 242–44]. Perry [2003: 1292] argues that the temporal account states a necessary, but no sufficient, condition for harm. Foddy [2014] defends a hedonist (and therefore not axiologically neutral) version of the temporal account.

⁷ Accounts along these lines have been defended by Feinberg [1984], Norcross [2005], Kloksiem [2012], and Purshouse [2016]. I thank Anna Folland for helpful suggestions on the formulation of the account.

⁸ Bradley [2012: 394-95] lists four additional desiderata that, however, will not feature in my discussion: prudential and normative importance of harm, amorality, and ontological neutrality.

My focus here is on *pro tanto* harm, rather than overall harm. The football player who twists her ankle while scoring an important goal might be overall better off for having scored (and receiving bonus pay), but nonetheless the twisted ankle is a *pro tanto* harm. In the following, I understand ‘harms’ and ‘benefits’ in the *pro tanto* sense, unless specified otherwise.⁹

2. Comparative Harm

According to comparative accounts, to find out whether an agent suffers harm, one compares the agent’s welfare in her actual state with the agent’s welfare in some comparison state. The most prominent comparative account is the counterfactual comparative account:

(Counterfactual Comparative) Agent suffers a harm if and only if an event occurs such that Agent would have been better off in the absence of that event.

The counterfactual comparative account faces several difficulties. To begin with, it cannot distinguish harms from failures to benefit. In a case given by Purves [2019: 2634], Batman intends to give golf clubs to Robin, but then tries them out and decides to keep them.¹⁰ The counterfactual comparative account says that Robin suffers harm in this case: Robin is worse off than he would have been, had Batman not tried out the clubs.

Moreover, the counterfactual comparative account does not recognize harms in cases where there is no event such that Agent would have been worse off in the absence of this

⁹ *Pro tanto* harms can be understood as harms relative to different *aspects of welfare*, as argued by Hanser [2008: 424], who speaks of harms ‘in a respect’. They can also be understood as different *contributors to welfare*. A view along these lines is Kagan’s ‘local’ interpretation of doing harm that is concerned with ‘the various individual elements that contribute to changes in well-being’ [Kagan 1998: 87]. For the purpose of this paper, I remain neutral on the question how we should individuate *pro tanto* harms. *Pro tanto* harms and overall harms are usually taken to be interdefinable: ‘an event is overall harmful to someone iff its *pro tanto* harms to that person outweigh its *pro tanto* benefits to that person’ [Bradley 2012: 393-94].

¹⁰ Purves’s case is a variant of a case given by Bradley [2012: 397].

event.¹¹ One example are non-identity cases in which an action that causes someone's existence also seems to harm them [Parfit 1984: ch. 16]. Another example are pre-emption¹² cases:

(Pre-emption) Al breaks Tim's nose. There is no event such that, had this event not occurred, Tim's nose would have remained intact (had Al not been there, then someone else would have broken Tim's nose instead).¹³

According to counterfactual comparativists, Tim does not suffer harm. This is implausible. It seems clear that the broken nose constitutes a harm.

I suggest that the counterfactual comparative account seems plausible because it captures intuitions about the moral relevance of actions: in other words, it captures intuitions about *harming*. The counterfactual comparative account identifies an action that makes a difference to someone's harmed state, and thereby explains why the agent is responsible for this state. Where an agent does not make a difference to another person's harmed state, it is more difficult to determine the agent's responsibility.

However, the counterfactual comparative account is much less plausible as an account of *harm*. Whether an agent suffers harm does not depend on counterfactual dependence relations between events and outcomes. To illustrate, imagine two cases in which Anne breaks her nose. In the first case, whether Anne breaks her nose counterfactually depends upon Bert's action. In the second case, Anne's nose will break whatever happens. There

¹¹ Another problem for counterfactual comparativists arises when it is indeterminate what would have happened had an event not occurred [Gardner 2017: 77].

¹² For responses to the pre-emption problem for counterfactual comparative accounts, see Feit [2015: 371], Boonin [2014: 62], Hanna [2016: 9-18], Kloocksiem [2012: 295-296]. For criticism of these responses, see Johansson and Risberg [2019].

¹³ This case differs from standard pre-emption cases, which merely assume that Tim's nose breaks regardless of Al's action. I make the stronger assumption that Tim's nose breaks in every possible world (at least, in every closest possible world to the events leading to Tim's injury). For example, imagine that in all possible worlds in which Tim exists, had Al not broken Tim's nose, then Villain would have broken Tim's nose, and had Villain not done so, someone else would have. Then, Tim does not suffer harm according to the counterfactual account. This is implausible. Thanks to an anonymous referee for pressing me on this.

might be a difference between these cases regarding whether Bert harms Anne. However, there is no difference between the two cases regarding whether Anne suffers harm. It is not plausible to say that whether the broken nose constitutes a harm depends on whether things could have turned out differently for Anne, or whether Bert's action brings about the broken nose. Anne is in a harmed state, regardless of whether there was a theoretical possibility of avoiding the broken nose.

In sum, the counterfactual comparative account faces objections: it does not recognize harm where it is present (as in pre-emption cases) and it does posit harm where it is not present (as in mere failures to receive benefits). I further suggested that the counterfactual comparative account explains intuitions about harming, but that it is less suited as an account about harm.

Given these worries, it is worth asking whether other accounts of harm are more plausible. Recall that the temporal comparative view replaces the counterfactual with a temporal comparison state, which describes the agent's state before suffering harm:

(Temporal Comparative) Agent suffers harm if and only if Agent is worse off than before.

The temporal account is rarely defended as a serious alternative to the counterfactual account. This is because it appears that it can easily be dismissed with reference to cases such as the following:

(Alleviate Suffering) The health of a terminally ill patient is rapidly deteriorating.

Doctor can alleviate her suffering by administering a drug, but she is unable to stop or

reverse the deterioration. Doctor administers the drug, and Patient's health deteriorates at a slightly slower pace.¹⁴

The temporal account seems to imply that Doctor harms Patient, because Patient is worse off after receiving the drug than before. This is implausible.

However, I deny that the temporal account of harm has this implication. It merely implies that Patient suffers harm (admittedly, less severe harm than Patient would otherwise have suffered). This seems right. What is implausible is not that Patient suffers harm. What is implausible is that Doctor's action counts as harming. But the temporal account of harm does not say anything about whether Doctor's action constitutes harming.

Another objection to the temporal account is that it cannot identify all harms:

(Delayed Recovery) Patient is about to recover when Doctor administers a drug that delays Patient's recovery.¹⁵

It might seem that the temporal account implies that Patient does not suffer harm, since Patient is just as badly off as before. However, Patient suffers harm in virtue of being worse off than Patient was before falling ill and needing treatment.¹⁶ Moreover, when Doctor prolongs Patient's suffering, Doctor causes Patient to remain in a harmful state for longer. Because severity of harm increases with the duration of suffering, Doctor makes the harm that Patient suffers more severe. But Doctor's action does not make a difference as to whether Patient suffers harm. (If Patient not only suffers harm, but Doctor also harms Patient, and the temporal comparative account of harming cannot explain why, then this shows that it is a worse account of harming than of harm.)

¹⁴ I adapted this case from Norcross's 'Doctor' case [Norcross 2005: 149].

¹⁵ Similar cases are given by Holtug [2002: 368], Petersen [2014: 204] and Hanser [2008: 429].

¹⁶ Thomson [2011: 444–45] suggests another response on behalf of the temporal comparativist: Patient might be worse off after receiving the drug in virtue of having worse recovery prospects (see also Velleman [2008: 243]). However, see Rabenberg [2014: 18] for objections to the prospect response.

Unfortunately, there is a case that is more deeply troubling for the temporal comparativist:

(Bad Start in Life) Patient was born with a painful condition and so, has never been better off.¹⁷

No matter how badly off Patient is, if Patient has always been this badly off, the temporal comparativist says that Patient is not in a harmed state. This seems implausible.¹⁸

3. Non-Comparative Harm

The main competitor of comparative accounts is the non-comparative account of harm:

(Non-Comparative) Agent suffers harm if and only if Agent is badly off.

On this account, a person is harmed if and only if she is intrinsically badly off. The problem with non-comparative harm is that there are cases in which it seems intuitively clear that a person suffers harm without being in an intrinsically bad state. Consider

(IQ) A professor with a very high IQ takes a drug that reduces her IQ by a few points.¹⁹

It seems that the professor suffers harm. However, by stipulation, the professor is not badly off: she is still very intelligent.

In response, non-comparativists might bite the bullet and accept that the professor is not suffering harm. Shiffrin seems to suggest this line of argument when she says, in her

¹⁷ Similar cases are mentioned in the literature as supposedly fatal objections to the temporal comparative account (Thomson [2011: 445-446], Holtug [2002: 369], Shiffrin [2012: 370]). See Velleman [2008: 244] for a bullet-biting response.

¹⁸ More generally, it might seem intuitively plausible that when someone is intrinsically badly off, this person is thereby in a harmed state, regardless of their previous welfare. The hybrid account accommodates this intuition.

¹⁹ I adapted this case from Hanser's 'Nobel Prize Winner' case [Hanser 2008: 432].

discussion of comparative accounts, that ‘comparative accounts . . . identify as harm cases in which one merely loses or fails to receive a tremendous benefit. A billionaire’s accidental loss of a thousand dollars will be said to be a harm to him, assuming he has a stake in his stockpile’ [Shiffrin 2012: 371].

However, biting the bullet comes at the price of revising the ordinary use of the notion of harm. It seems that the professor is in a harmed state, having lost some of her cognitive powers. Non-comparativists who bite the bullet must, however, insist that this is incorrect: the professor has merely lost a benefit.

Moreover, in everyday usage we distinguish between welfare losses and mere failures to receive benefits. Welfare losses are morally significant in a way in which mere failures to receive benefits are not. For example, welfare losses, but not failures to receive benefits, generally ground claims to compensation. This significance is explained by the intuitively plausible claim that someone who loses welfare suffers harm, but someone who merely fails to receive a benefit does not suffer harm. Non-comparativists who bite the bullet cannot accept this claim. If the professor’s loss is not a harm, but merely a lost benefit, then entitlements to compensation must be defended on non-harm-based grounds.

4. The Hybrid Account: The Perfect Match?

We have reached a stalemate. Cases like (Bad Start in Life) show that one can be in a harmed state without being comparatively worse off. Cases like (IQ) show that one can be in a harmed state without being non-comparatively badly off.

I think that the solution is that we should adopt a hybrid account, combining comparative and non-comparative accounts. In general, a hybrid account says that an agent suffers harm if and only if an agent’s welfare is either lower than a non-comparative threshold

of harm, or lower than a comparative baseline. More formally, let A be an agent with welfare level W. Then, for a non-comparative threshold T, and a comparative baseline C, the basic idea is this:

(Hybrid Formula) A suffers harm if and only if either $(W < T)$ or $(W < C)$.

Fulfilling either component of the hybrid formula is sufficient for harm and fulfilling at least one is necessary. In the following, I will defend a version of the hybrid account, which combines a temporal comparative account of well-being and a non-comparative account of ill-being. By an agent's overall *well-being*, I mean the sum of elements that make the agent's life go well, and by an agent's overall *ill-being*²⁰ I mean the sum of elements that make an agent's life go badly. I will use the term *welfare* to refer to how an agent fares overall (which could be well or badly). I assume that someone who is neither in a good state nor in a bad state has a neutral welfare level. Well-being and ill-being are defined with reference to this neutral welfare level. A level of welfare that is above the neutral level constitutes well-being, and a level of welfare that is below the neutral level constitutes ill-being.

We have seen in the previous discussion that the temporal comparative account of harm fails because it cannot identify instances of ill-being that people are born into as harms, however, it can identify loss of well-being as harm. The non-comparative account fails because it cannot identify loss of well-being as harm, however, it can identify instances of ill-being that people are born into as harms. As I will argue, the combination of these accounts is extensionally adequate:

(Hybrid) A suffers harm if and only if either (i) A suffers ill-being or (ii) A's well-being is lower than it was before.

²⁰ For discussion on the concept of ill-being, see Kagan [2014].

The temporal comparative account here is slightly altered: rather than applying to welfare, it applies to well-being only. As I explain in the next section, this allows the hybrid account to be presented as a unified account, rather than an ensemble of unrelated disjuncts.²¹

4.1 Is the Hybrid Account Unified Enough?

The hybrid account might seem dubiously ad hoc. Even if it gets the cases right, it seems to tell two different stories about the nature of harm, rather than one story about what the unified core of harm is. In other words, an objection to the hybrid account is that it cannot explain what makes cases such as (IQ) and (Bad Start in Life) two instances of the same phenomenon, rather than two distinct kinds of cases.

It might be tempting for the proponent of the hybrid account to argue that if a hybrid account fulfils the desideratum of extensional adequacy, then perhaps giving up the desideratum of unity is a price worth paying.

However, I do not think proponents of the hybrid account have to give up unity. Agents who suffer either comparative or non-comparative harm have one thing in common: their welfare is adversely affected. Agents who enter a state of being comparatively worse off than before, or being non-comparatively badly off, suffer an adverse effect on their actual welfare. In contrast, the actual welfare of someone who merely fails to receive a benefit is not

²¹ Meyer's 'disjunctive notion of harm' [Meyer 2016: sec. 3.4] is a different version of the hybrid account. Roughly, the disjunctive notion combines a qualified version of the causal account of harming with a counterfactual comparative account. However, the disjunctive notion faces counterexamples. For example, consider the case of the lorry driver who, despite taking all precautions, runs over a child who jumps into the road. Had he not run over the child, the driver behind him would have done so. Surely, the child suffers harm and the lorry driver has harmed the child (although the driver may not be blameworthy). Meyer's disjunctive account does not explain this judgement. The first disjunct does not explain it, because, by stipulation, the agent could not avoid causing harm and did minimize the harm (as far as possible). The second disjunct does not explain it, because it is not true that the child would have been better off, had the driver not run over the child, since in that case the second lorry would have run over the child.

adversely affected. The welfare of such a person remains positive. What changes is merely how the person's actual welfare compares to hypothetical baselines in possible worlds.

At this point, we can see why it is useful to restrict the scope of the comparative account to well-being, and the scope of the non-comparative account to ill-being. It allows us to see the difference between the two ways in which an agent can suffer harm according to the hybrid account: by losing well-being, and by gaining (or continuing to occupy) ill-being. In other words, an agent can be in a harmed state in virtue of being non-comparatively badly off, or in virtue of being worse off than before. In both cases, the agent's actual welfare is adversely affected.

4.2 Is the Hybrid Account Really Extensionally Adequate?

The hybrid account is extensionally adequate if and only if it recognizes all cases of harm as such and does not find harm where there is none. I argue that the hybrid account correctly identifies all cases of harm. The hybrid account's first clause can explain why Tim suffers harm in (Pre-emption). Tim is non-comparatively badly off: Tim is in a state of ill-being. The hybrid account's second clause can explain why the professor suffers harm in (IQ). It is because the professor is comparatively worse off: the professor's well-being is lower than it was previously.

However, there are cases in which the hybrid account seemingly fails to recognize harm where it exists. These cases are preventions of non-comparative benefits. Consider

(Lottery Ticket) Ann prevents Bob from buying a lottery ticket. If Bob had bought the ticket, he would have won.

Bob is not badly off, and Bob is not worse off than he was before. According to the hybrid account, Bob does not suffer harm. However, it is not clear that Bob does not suffer harm.

After all, if Ann had not prevented Bob from buying a ticket, Bob would now be a rich man. So, it might seem that Ann has harmed Bob. However, Ann cannot have harmed Bob if Bob is not in a harmed state. So, either Ann harms Bob, in which case the hybrid account in the version that I am defending here is incorrect, or Ann does not harm Bob and intuitions to the contrary are mistaken.

I think that the hybrid account is correct in saying that Bob does not suffer harm. Rather, Bob does not receive a benefit.²² According to the hybrid account, whether Bob is in a harmed or benefitted state is independent of whether he could have been in a better or worse state in counterfactual scenarios. Since Bob does not suffer harm, Ann does not harm Bob. She prevents him from receiving a benefit.

It seems intuitively correct that Ann's behaviour is prima facie objectionable.²³ However, the hybrid account need not deny this. Ann actively prevents Bob from receiving a benefit. Preventing Bob from receiving the benefit seems morally problematic in a way that merely failing to benefit Bob is not.

The objection can be pushed further. In (Potion), the causal sequence leading to the benefit has already started:

(Potion) Cleo gives Dora a potion, which will increase Dora's athletic ability. Eric gives Dora another potion, which prevents the first potion from taking effect.²⁴

²² On some views of welfare, possessing a winning lottery ticket might constitute a benefit even before the draw. On such views, Bob would be in a harmed condition when the ticket is replaced after he bought it, but not when he is prevented from buying it.

²³ When we change the case such that Ann's behaviour does not seem objectionable, the case no longer seems to involve harm. For example, imagine that Bob regularly spends money on lotteries. His chances of winning are very low. Knowing this, Ann persuades Bob to pay the money into his pension instead. Had Bob played, he would have won. This case does not seem to involve harm, and yet Bob is in exactly the same state.

²⁴ This case has been suggested to me by an anonymous referee. I am grateful to the referee for the case and for raising this objection.

It might seem that Eric harms Dora by giving her the second potion. Is Dora in a harmed state, according to the hybrid account? I suggest that this depends on the account of well-being that is coupled with the hybrid account. Some accounts of well-being might imply that Dora's well-being increases when she drinks the first potion. (For example, consider the view that extraordinary capability or talent for future excellence can contribute to present welfare.) According to this view, the hybrid account implies that Dora is in a harmed state after receiving the second potion.

Other accounts of well-being might imply that Dora's welfare does not increase when she drinks the first potion. On such views, the hybrid account implies that Dora is not in a harmed state in (Potion). I think that this implication is plausible. Dora does not suffer harm when her athletic performance remains stable. (Imagine a case in which Dora takes dietary supplements for muscle growth that unfortunately neutralize each other. Dora fails to receive a benefit, but she is arguably not in a harmed state.

What is implausible is the view that Eric's behaviour is not objectionable. However, similar to Ann in (Lottery Ticket), the intuition that Eric's behaviour is objectionable might be explained by the fact that Eric actively prevents Dora from receiving a benefit.

I conclude that whether Dora is in a harmed state in (Potion) depends on the account of well-being that is coupled with the hybrid account. Moreover, the proponent of the hybrid account can explain why Eric's behaviour is objectionable even if Dora does not suffer harm.

A further worry for the hybrid account arises from cases like the following:

(Toothache) Cinderella suffers from a very painful toothache. The nice fairy gives her a potion that reduces the pain significantly, but does not reduce it to zero.²⁵

²⁵ I slightly adapted this from a case generously given to me by an anonymous referee.

It seems like the nice fairy benefits Cinderella. But the hybrid account says that Cinderella is still in a harmed state after drinking the potion. So, it might seem that the hybrid account finds harm where there is none. However, Cinderella is in a harmed state (due to the remaining pain) *and* in a benefitted state (due to having less ill-being than previously). So, the hybrid account correctly identifies harm in this case. More precisely, the hybrid account correctly says that Cinderella is both in a non-comparatively harmed state and in a comparatively benefitted state.

What these cases illustrate is the need for an account not only of harm, but also of benefits. In the next section, I argue that the hybrid view provides such an account.

Before I do so, let me first point out that the hybrid account can also account for preventive harms:

(Ambulance) Ann stops the ambulance that is on its way to Bob, who is in pain.

In this case, the hybrid account implies that Bob suffers harm. He is clearly badly off in virtue of suffering ill-being. Since Bob suffers pain for longer (while he waits for assistance), it is possible (and in this case plausible) that Ann has harmed Bob. In contrast, suppose that Ann stops the delivery driver who is about to deliver Bob's Christmas present. In this case, Bob is not in a harmed state. He merely fails to receive a benefit.

4.3 Can the Hybrid Account Explain Benefits?

There is a hybrid account of benefit that is analogous to the hybrid account of harm:

(Hybrid-Harm) A suffers harm if and only if either (i) A suffers ill-being or (ii) A's well-being is lower than it was before.

(Hybrid-Benefit) A enjoys a benefit if and only if either (i) A enjoys well-being or (ii) A's ill-being is lower than it was before.

In (Toothache), the first condition of (Hybrid-Harm) and the second condition of (Hybrid-Benefit) are fulfilled. Cinderella enjoys a benefit (the tooth hurts less than it did), but she also suffers harm (the tooth is painful). In contrast, in (IQ), the second condition of (Hybrid-Harm) and the first condition of (Hybrid-Benefit) are fulfilled. The professor suffers harm (she loses cognitive abilities), but the professor also enjoys a benefit (she is very intelligent).

In these cases, the same state of affairs constitutes a harm and a benefit. This might seem odd at first sight. However, it should not strike us as odd once we remember the nature of the hybrid account: the hybrid account identifies two ways in which someone's welfare can be affected. In the cases just described, the agents' welfare is affected in different ways. This is why these states constitute a harm in one sense and a benefit in another sense. This corresponds to intuitions about harm, which pull in both comparative and non-comparative directions.

To explain this further, it is helpful to consider Hanser's distinction between etiological and non-etiological conditions. According to Hanser [2019: 859], etiological conditions imply something about the history leading to the condition, whereas non-etiological conditions are independent of the history leading to the condition. Hanser argues that being in a harmed state refers to an etiological condition only [see Hanser 2019: 858–60].

I agree with Hanser that we often speak of harmed and benefitted states as saying something about how agents got in that state, or the state they were in previously. The temporal comparative account of harms and benefits captures this intuition. However, *pace* Hanser, I think that by saying that someone is in a harmed condition, we sometimes want to say that they are badly off, where it does not matter how they came to be in this condition.

This is the intuition behind (Bad Start in Life). Here, harms and benefits say something about the intrinsic badness or goodness of a state for an agent. The non-comparative account captures this intuition.

A state can be good for an agent even if she was in a better state previously, as in (IQ). A state can be bad for an agent who was in an even worse state previously, as in (Toothache). In these cases, comparative and non-comparative intuitions pull in different directions.

In sum, I suggest that there is nothing odd about the idea that the same condition can be a harmed condition for someone in the non-comparative sense (relative to a non-comparative threshold), but a benefitted condition in the temporal comparative sense (relative to a temporal baseline), or vice versa. The hybrid account of harm and benefit can explain these complicated cases, and how they differ from simpler cases in which a state constitutes a non-comparative harm (or benefit) and also a comparative benefit (or harm). This also suggests that there is no straightforward answer to the question whether the agent in cases like (IQ) or (Toothache) is in a harmed or benefitted condition *overall*.²⁶

4.4 Can the Hybrid Account Explain the Severity of Harm?

If an account of harm should inform decision making, it needs to enable us to measure the severity of harm and benefit. In this section, I explore how defenders of the hybrid account

²⁶ It might seem tempting to think that it is permissible to add up comparative and non-comparative harms and benefits. However, this is implausible. Consider a case in which Ann's welfare level is 5, and Bob's is -5. Assume that Ann's previous welfare level was 15, whereas Bob's has always been at -5. If it were permissible to add up all pro tanto harms and benefits, then we could reason as follows. Ann has lost 10 units of well-being, and has 5 units of well-being left. So, she suffers a harm of 10 and a benefit of 5, which gives her an overall harm of -5. It therefore seems that Ann and Bob are in an equally severe harmed state. This seems very implausible. In fact, it seems important to distinguish the balance of comparative harms and benefits and the balance of non-comparative harms and benefits. I tentatively suggest that, to calculate an agent's *overall welfare*, we can add up non-comparative pro tanto harms and benefits. To calculate the *overall magnitude of changes* to an agent's welfare, we can add up pro tanto temporal comparative harms and benefits. Importantly, we cannot determine whether an agent is *in a harmed state overall* by adding up comparative and non-comparative pro tanto harms. Instead, we need to look at these dimensions separately. An agent can be harmed in a comparative sense, yet not in a non-comparative sense, and vice versa.

should think about the severity of harm. My aim is not to conclusively defend an account of the severity of harm, but to draw attention to some questions that arise in formulating such an account.

The following claim seems plausible:

(Harm Severity) The severity of a harm is proportional to the extent to which Agent's welfare is adversely affected.

Then, the non-comparative component of the hybrid account implies that other things being equal, harm is more severe the more ill-being the agent suffers, and the comparative component of the hybrid account implies that harm is more severe the more well-being the agent has lost.

We can make an analogous claim for benefits:

(Benefit Size) The size of a benefit is proportional to the extent to which Agent's welfare is positively affected.

One way to measure the severity of temporal comparative harms is simply to measure the units of welfare lost. However, measuring the severity of harm in this way fails to take into account the duration of harm. It implies that a temporary loss of well-being is just as severe as a permanent loss of the same amount of well-being. This is implausible. The severity of temporal comparative harm, then, does not only depend on the magnitude of the loss of well-being, but also on its duration. A very simple way to calculate the severity of harm based on magnitude of loss and duration would be to multiply the units of well-being lost by the time during which the well-being remains unrecovered. For example, imagine that A's well-being drops from 100 to 50 for 10 years. Then, the severity of harm is 500: the amount of welfare units lost (50) multiplied with the duration of the loss in years (10). Now, this way of calculating the severity of comparative harm is illustrative but overly simple, and I do not wish to

claim that magnitude and duration of welfare losses are the *only* factors that determine the severity of temporal comparative harms. I merely suggest that these two factors ought, minimally, to be considered.²⁷

Now, consider the following case, suggested to me by an anonymous referee:

(Double Effect) Dora comes into existence with a welfare level of 100. Unless someone intervenes, Dora will occupy this level for 100 years and then die. Eric gives Dora a potion that has two effects. First, the potion immediately lowers Dora's welfare level to 0. Second, the potion raises Dora's welfare by 1 every year, until Dora dies at level 100.

It seems clear that Dora is in a harmed state. However, the hybrid account seems to imply that Dora is not in harmed state overall. Adding up the temporal comparative harms that Dora suffers relative to the initial situation over the years, we get $100+99+98+\dots+1+0=5050$. Adding up the non-comparative benefits that Dora enjoys over the years, we get $0+1+2+\dots+100=5050$. It seems that the hybrid account implies that Dora is not in a harmed state overall, since $5050-5050=0$.²⁸

The objection relies on the assumption that the extent to which an individual is 'overall' in a harmed state can be determined by adding up comparative and non-comparative harms and benefits. However, as I argued in the last section, we should reject this assumption.

The assumption that we can determine whether an agent is in an 'overall' harmed state by adding up comparative harms and non-comparative benefits is implausible. Consider a variant of (Double Effect), in which Dora's welfare level is 200, then drops to 100 and slowly rises back to 200. The comparative harm that Dora suffers seems equal to the comparative harm that Dora suffers in the original case. However, the non-comparative benefit that Dora enjoys is significantly higher than in the original case. So, the assumption that we can add up

²⁷ Thanks to an anonymous referee for pressing me to clarify this.

²⁸ This line of argument was suggested by the anonymous referee who generously provided the case.

these harms and benefits seems to have the implication that Dora is in a benefitted state overall. This is implausible.

The point that we cannot simply add up comparative and non-comparative harm and benefits makes sense when we reflect on the nature of comparative and non-comparative harms. Comparative and non-comparative harms are relative to different baselines. Dora is, at the same time, in a comparatively harmed state (relative to a temporal baseline) and a non-comparatively benefitted state (relative to a non-comparative baseline). We should not expect to derive useful results from adding up harms and benefits that are relative to different baselines.²⁹

If it is not clear whether Dora is in a harmed state overall, it might seem to follow that there is no answer to the question whether the potion harms Dora. However, this does not follow. Dora clearly is in a comparatively harmed state. Eric's behaviour is plausibly related in a harming way to the comparative harm Dora suffers. The hybrid account does not provide an account of this harming relation. However, the hybrid account is compatible with the following views: the view that the potion harms Dora, the view that the potion harms Dora more than it benefits her, and even the view that the potion does not benefit Dora at all. (The view that the potion does not benefit Dora can be based on counterfactual considerations, since Dora would have enjoyed less comparative harms, but not more non-comparative benefits, had she not drunk the potion. Remember that the hybrid account does not make any claims about harming.)³⁰

²⁹ See footnote 26 for an example that illustrates this point.

³⁰ Another potential response to (Double Effect) might build on the claim that duration and intensity of harm are not the only factors that can influence the severity of harm. Perhaps the harm is more severe than the benefit in this case because it occurs all at once. The sudden loss weighs more heavily than the slow increase in well-being. However, we can imagine a case in which Dora receives the benefits all at once (her well-being stays at 0 before it jumps back to 100). In this case, it still seems that Dora is in a harmed state.

I conclude that the objection to the temporal comparative account posed by (Double Effect) fails. The hybrid account does not imply that Dora does not suffer harm in this case. It also does not imply that Eric does not harm Dora.

The objection can be pushed further still. Consider

(Single Effect) Freddie's well-being level is 100. Ginger gives Freddie a potion that causes Freddie's well-being level to drop to 1 and to remain there until Freddie dies.

(Pre-empted Single Effect) Freddie's well-being level is 100. Ginger gives Freddie a potion that causes Freddie's well-being to drop to 1 and to remain there until Freddie dies. Had Freddie not drunk the potion, Freddie's well-being level would have dropped to 1 naturally.

The argument that I gave above implies that Freddie suffers equivalent temporal comparative harm in (Single Effect) and in (Pre-empted Single Effect). This is because in both cases, Freddie's well-being drops to 1 and remains there until she dies. The severity of harm remains the same. This might seem implausible. After all, in (Pre-empted Single Effect), the potion makes a much smaller difference to Freddie's lifetime welfare, since Freddie's welfare would have dropped naturally in any case.

Here, too, it is worth keeping in mind the distinction between harm and harming. It seems plausible that Freddie suffers equivalent harm in both cases, since she experiences a loss in well-being. What plausibly changes is the extent to which Ginger's potion is responsible for the harm. The hybrid account, however, is only concerned with harm. The hybrid view is compatible with the view that Ginger harms Freddie to a greater extent in (Pre-empted Single Effect) than in (Single Effect).³¹

³¹ Again, counterfactual considerations might explain this view. In (Pre-empted Single Effect), had Freddie not drunk the potion, she would have been just as badly off. So, the potion does not make much difference to Freddie's welfare in (Pre-empted Single Effect), but it does in (Single Effect). This plausibly changes the extent to which Ginger is responsible for the harm Freddie suffers.

5. Conclusion

I argued that the best account of harm combines comparative and non-comparative elements, and proposed that we should adopt a hybrid account of harm. A version of the hybrid account that includes a temporal comparative well-being baseline and a non-comparative ill-being baseline is extensionally adequate and provides the resources to be presented as a unified account of harm. Further, I tentatively argued that the hybrid account of harm gives rise to a corresponding hybrid account of benefit, and that magnitude and duration influence the severity of harm. A general lesson that can be drawn from my discussion is that the question of what constitutes harm and the question of what constitutes harming are importantly distinct. An analysis of harm should make this distinction explicit. Moreover, distinguishing different kinds of harm seems to be a promising way to ensure extensional adequacy. Hybrid accounts can form a framework for different kinds of harm and facilitate the analysis of the moral significance of different kinds of harm.

Acknowledgements

I am very grateful to three anonymous referees for extremely detailed and thoughtful comments. For excellent comments and helpful discussion, I owe a great debt of gratitude to Fiona Woollard, Brad Hooker, Matthew Hanser, Brian McElwee, and Anna Folland. I thank participants in the Higher Seminar in Practical Philosophy at Uppsala University and the research group at the Chair for International Relations at the Technical University of Munich for helpful discussion of previous versions of this paper.

REFERENCES

- Boonin, David 2014. *The Non-Identity Problem and the Ethics of Future People*, New York: Oxford University Press.
- Bradley, Ben 2012. Doing Away with Harm, *Philosophy and Phenomenological Research* 85/2: 390–412.

- Feinberg, Joel 1984. *Harm to Others: The Moral Limits of the Criminal Law Volume One*, New York: Oxford University Press.
- Feit, Neil 2015. Plural Harm, *Philosophy and Phenomenological Research* 90/2: 361–88.
- Foddy, Bennett 2014. In Defense of a Temporal Account of Harm and Benefit, *American Philosophical Quarterly* 51/2: 155–65.
- Gardner, Molly 2017. On the Strength of the Reason Against Harming, *Journal of Moral Philosophy* 14/1: 73–87.
- Hanna, Nathan 2016. Harm: Omission, Preemption, Freedom, *Philosophy and Phenomenological Research* 93/2: 251–73.
- Hanser, Matthew 1990. Harming Future People, *Philosophy & Public Affairs*, 19/1: 47–70.
- Hanser, Matthew 2008. The Metaphysics of Harm, *Philosophy and Phenomenological Research* 77/2: 421–50.
- Hanser, Matthew 2019. Understanding Harm and Its Moral Significance, *Ethical Theory and Moral Practice* 22/4: 853–70.
- Harman, Elizabeth 2004. Can We Harm and Benefit in Creating?, *Philosophical Perspectives* 18/1: 89–113.
- Harman, Elizabeth 2009. Harming as Causing Harm, in *Harming Future Persons: Ethics, Genetics and the Nonidentity Problem*, ed. Melinda A. Roberts and David T. Wasserman, Dordrecht: Springer: 137–54.
- Holtug, Nils 2002. The Harm Principle, *Ethical Theory and Moral Practice* 5/4: 357–89.
- Johansson, Jens, and Olle Risberg 2019. The Preemption Problem, *Philosophical Studies* 176/2: 351–65.
- Kagan, Shelly 1998. *Normative Ethics*, New York: Routledge.
- Kagan, Shelly 2014. An Introduction to Ill-Being, in *Oxford Studies in Normative Ethics Volume 4*, ed. Mark Timmons, Oxford: Oxford University Press: 261–88.
- Klocksiesm, Justin 2012. A Defense of the Counterfactual Comparative Account of Harm, *American Philosophical Quarterly* 49/4: 285–300.
- McMahan, Jeff 2013. Causing People to Exist and Saving People’s Lives, *The Journal of Ethics* 17/1-2: 5–35.
- Meyer, Lukas H 2016. Intergenerational Justice, *The Stanford Encyclopedia of Philosophy (Summer 2016 Edition)*, ed. Edward N. Zalta, URL = <https://plato.stanford.edu/archives/sum2016/entries/justice-intergenerational/>.
- Norcross, Alastair 2005. Harming in Context, *Philosophical Studies* 123/1–2: 149–73.
- Parfit, Derek 1984. *Reasons and Persons*, New York: Oxford University Press.
- Perry, Stephen 2003. Harm, History, and Counterfactuals, *San Diego Law Review* 40/4: 1283–1313.
- Petersen, Thomas Søbirk 2014. Being Worse off: But in Comparison with What? On the Baseline Problem of Harm and the Harm Principle, *Res Publica* 20/2: 199–214.
- Purshouse, Craig 2016. A Defence of the Counterfactual Account of Harm, *Bioethics* 30/4: 251–59.
- Purves, Duncan 2019. Harming as Making Worse Off, *Philosophical Studies* 176/10: 2629–56.
- Rabenberg, Michael 2014. Harm, *Journal of Ethics and Social Philosophy* 8/3: 1–32.
- Rivera-lópez, Eduardo 2009. Individual Procreative Responsibility and the Non-Identity Problem, *Pacific Philosophical Quarterly* 90/3: 336–63.
- Shiffrin, Seana Valentine 1999. Wrongful Life, Procreative Responsibility, and the Significance of Harm, *Legal Theory* 5/2: 117–48.
- Shiffrin, Seana Valentine 2012. Harm and Its Moral Significance, *Legal Theory* 18/3: 357–98.
- Thomson, Judith Jarvis 2011. More on the Metaphysics of Harm, *Philosophy and Phenomenological Research* 82/2: 436–58.

Velleman, James David 2008. The Identity Problem, *Philosophy & Public Affairs* 36/3: 221–44.

Woollard, Fiona 2012. Have We Solved the Non-Identity Problem?, *Ethical Theory and Moral Practice* 15/5: 677–90.

Funding Information

This research was funded by an Arts and Humanities Research Council Grant (AH/L/503939/1).

Institutional Affiliation

University of Southampton, Technical University of Munich