



Explaining Harm

Eli Pitcovski¹

Accepted: 6 December 2022

© The Author(s), under exclusive licence to Springer Nature B.V. 2022

Abstract

What determines the degree to which some event harms a subject? According to the counterfactual comparative account, an event is harmful for a subject to the extent that she would have been overall better off if it had not occurred. Unlike the causation based account, this view nicely accounts for deprivational harms, including the harm of death, and for cases in which events constitute a harm rather than causing it. However, I argue, it ultimately fails, since not every intrinsically bad state that is counterfactually dependent on an event contributes to its degree of harm. So while the causation based account is too restrictive, the counterfactual comparative view is not restrictive enough. In light of this, I suggest an alternative, explanation based account of overall harm, according to which the degree to which some event is harmful for a subject is determined by the degree to which (crudely) the states explained by it are overall more intrinsically bad than intrinsically good for her.

Keywords Causal explanation · Counterfactuals · Wellbeing · The counterfactual comparative account of harm · The harm of death

1 Introduction

The notions of harm and benefit play a key role in the evaluation of events we are subjected to, and of alternative courses of actions we plan to undertake. Nevertheless, they are difficult to account for. In what follows I will be interested in what makes an event overall harmful (rather than *pro-tanto* harmful), and in what determines the degree to which an event is overall harmful.

The most straightforward attempt to account for overall harm is to associate the harm of an event with the intrinsic badness of what it causes (I will use ‘bad’ and ‘harmful’ interchangeably). For instance:

✉ Eli Pitcovski
elielipit@gmail.com

¹ Tel-Hai Academic College, 9977 road, 1220800 Qiryat Shemona, Israel

CAUSE: An event *E* is harmful for some subject *S* iff the total consequences of *E* are more intrinsically bad than intrinsically good for *S*. The degree¹ to which an event *E* is harmful for some subject *S* is the degree to which the total consequences of *E* are more intrinsically bad than intrinsically good for *S*.²

A familiar complaint about CAUSE is that it fails to account for the harm of death, and other cases of harm due to deprivation. Normally, death at the age of twenty is presumably harmful for its subject. But if, as I will assume throughout, a subject's death is simply their ceasing to exist, death brings about no intrinsically bad states for them. It is only harmful by virtue of the good states it deprives its subject of. Since CAUSE associates the harm of events strictly with the value of the states they cause, it cannot account for the harmfulness of death (or so goes the familiar complaint).

In light of this, the most prevalent approach to overall harm is the counterfactual comparative account, according to which an event is harmful for a subject to the extent that they are worse off in comparison to what would have happened if it did not take place. On a rough formulation,³ it is the view that:

COMPARE: An event *E* is harmful for some subject *S* iff the actual world (in which *E* occurs) is overall worse for *S* than the nearest possible world in which *E* does not occur. The degree to which an event *E* is harmful for *S* is the degree to which the actual world is worse for *S* than the nearest possible world in which *E* does not occur.

COMPARE accounts for deprivational harm, including the harm of death. If *S* dies of an accident at the age of twenty, the event of his death does not bring about states that are intrinsically bad for *S*. But if it were not for this particular death,⁴ *S* would have continued to lead a life that is, presumably, more intrinsically good than bad. More generally: the life *S* would have had if *S* had not died in those circumstances is better overall than the life he actually lives. The

¹ Not all adherents of causation-based accounts and of the counterfactual comparative account discussed below address degrees of harm. Given that my interest is an account that captures the degree of harm, I focus on formulations that do.

² See Shiffrin (2012) for this sort of causation based account.

³ Some counterfactual comparative accounts involve modifications of COMPARE in light of objections discussed in Sect. 5. (See: Purves 2019; Klockslem 2012; Feit 2015; Bradley 2004, pp. 18–21; 2009, 2.2, Hanna 2015). Moreover, some versions account for overall benefit on top of accounting for overall harm. For the sake of concise presentation of the main problem, it will be easiest to focus on COMPARE for now. For similar formulations, see Bradley (2009, p. 50), Broome (1999; 2004); Feldman (1991), p. 150).

⁴ How is a particular event of death individuated? In order for COMPARE to get off the ground, events must not be taken to be extremely modally fragile (with the stringent identity conditions). Otherwise: the nearest world in which some event of death does not take place will always be a world in which a very similar death takes place an extremely short period after it. This delivers a counterintuitive result COMPARE was supposed to avoid, viz. that individual deaths are essentially unharmed, since what is deprived by them is negligible. Bradley (2009, p. 56) allows for contexts in which 'particular death' is understood in a way that renders the designated event modally fragile, but emphasizes that such contexts are "odd" (ibid.), and that there are normal contexts in which 'particular death' is not taken as modally fragile.

intuition that the deprivation of an overall good life period is a harm is nicely captured by COMPARE. Rather than focusing on the actual states resulting from some particular event of death, COMPARE focuses on how much better a subject's life would have been if it weren't for this case of death.

COMPARE has other advantages over CAUSE. For instance, unlike CAUSE, COMPARE seems to account for constitutional harms: harms that are constituted by a certain event rather than caused by it. To use an example from Purves (2019, p. 2638): "Suppose that my mother desires that I not join a cult. If I join a cult, this frustrates her desire... But, it would misconstrue things to say that my joining a cult causes her desire to be frustrated... My joining a cult harms my mother by *constituting* her desire frustration." Hence, strictly speaking CAUSE fails to account for my mother's harm. COMPARE, on the other hand, correctly implies that my mother is harmed by my joining the cult: In the nearest possible world in which I do not join the cult, my mother's desire is not frustrated. We will come back to this point in Sect. (5), when considering a causation-based account that fares better than CAUSE in other respects. For the time being my focus will be on the supposed main advantage of COMPARE: accounting for deprivational harms.

In what follows I present a problem for COMPARE (Sect. 2), and after discussing it (Sect. 3), I provide the beginnings of another account, centering on the notion of explanation. I argue that this account overcomes the problem in question (Sect. 4) and that it is otherwise attractive (Sect. 5). On the suggested account an event is harmful for a subject to the extent that (crudely) the states explained by this event are altogether more intrinsically bad than intrinsically good for that subject.

2 A Problem for the Counterfactual Comparative Account

The problem with COMPARE I wish to focus on is that in some cases, rather than establishing the degree to which some event E is harmful, either it gives us instead the degree to which some events causally leading to E are harmful, or else it leads to absurdity. To illustrate the problem, consider the following case:

JOSEPH: Joseph's town is captured by a cruel and powerful army. The people of the town are given the following choice: either they execute one person and leave the body outside the walls by midnight, or the whole town will be subjected to a lifetime of torture. Since the commander who gave them the choice is reliable and has a reputation for carrying out his threats, the people of the town decide on a fair lottery that will determine which of them will be executed. By bad luck, Joseph is chosen. He is injected with a dose of poison strong enough to kill an elephant and then put in a cage outside the walls. The poison causes him pain that is worse than any pain he has ever experienced and by midnight, it causes his death. The rest of his town is not tortured.

In order to determine the degree to which Joseph's death is harmful for him according to COMPARE, we ought to compare the actual world to the nearest possible world in which this death fails to take place. What would that world look like? Here are two natural candidates:

OPTION ONE: The nearest world in which Joseph's death fails to take place is a world in which Joseph is not poisoned because someone else is chosen by the lottery.

OPTION TWO: The nearest world in which Joseph's death fails to take place is a world in which Joseph is poisoned but does not die as a result of it.

If we assume OPTION ONE, COMPARE comes close, but ultimately fails to establish the degree to which Joseph's death is harmful due to a familiar, often neglected problem with backtracking i.e. with the idea that the nearest world in which some event, E, fails to take place has a different past relative to the time of E. Comparing the actual world to the nearest possible world in which Joseph does not die does not give us the overall harm of his death. Instead, it gives us the overall harm of death together with some events preceding it (roughly, the overall harm of his being poisoned). After all, the actual world is worse for Joseph than an OPTION ONE world, in which he is not poisoned, partly because of the terrible pain he actually suffers before he dies. This pain has nothing to do with the (presumably essentially deprivational) harm of Joseph's death; nonetheless, COMPARE wrongly takes the absence of these moments of pain in the possible world in question into account when considering the harm of death. The absence of events preceding death in a world in which he is not poisoned positively contributes to the overall intrinsic value of Joseph's life in that world (compared to the actual world). That makes the comparison to that world misleading, given that we are interested in the value of Joseph's death, rather than the value of the process of his dying.⁵

OPTION TWO, on the other hand, has implausible consequences especially when coupled with COMPARE. If we compare the actual world in which Joseph dies of the poison to the nearest possible world in which Joseph is poisoned but does not die in result, Joseph's death would not seem very unfortunate; perhaps even beneficial for him overall. Keeping the rest of the setup intact, in that scenario, Joseph is tortured with the rest of the town for the rest of his life. For his survival means that the town has failed to present a body by midnight. In a normal context, this result is implausible. The assertion that Joseph's death was overall very good for him cannot (adequately) be assumed to be accommodated offhand.⁶ Of course, there may

⁵ There are perhaps conversational contexts in which we rightly associate the pain preceding death with the harms of a particular death. I will discuss such contexts below (in Sect. 5).

⁶ What I mean by saying that an assertion about the case in question is accommodated by a normal context is roughly that it meets the following (sufficient) condition: when asserted to an arbitrary individual that knows the nonlinguistic facts about the case but who cannot be presupposed to share any wide background knowledge or purposes, no communicative effort (e.g. no further explanations on behalf of the speaker or interpretive effort on behalf of the audience) is adequately assumed to be required for understanding what is said. Some things we might say, such as 'Joseph's death is highly beneficial for him' may be accommodated while requiring interpretational effort yielding a context shift.

be contexts in which Joseph's death should not be considered very bad for him; for instance in a context in which we take his poisoning as a given. (Such contexts are discussed in Sect. 5) Yet there is a perfectly normal context which accommodates the correct judgment that Joseph's premature death is terribly bad for him. Note that it is mainly the badness of his death that explains why it was such bad luck for Joseph to be chosen by the lottery.⁷

These considerations make OPTION ONE far more fitting. On any plausible account of counterfactuals,⁸ a world in which a fair lottery turns out differently is presumably much closer to the actual world than a world in which an injection strong enough to kill an elephant fails to kill Joseph. While change in the actual results of a lottery would require relatively small divergence from the actual world, the failure of the poison to kill Joseph would require either a great divergence in facts (some history that would make Joseph resistant to an enormous amount of poison), or a very big miracle.⁹ This gives rise to our general problem: when the nearest non-E worlds differ from actuality with respect to some events preceding E, and when these events are by themselves intrinsically harmful/beneficial, assessing the harm of E cannot be achieved by simple comparison provided by a theory of counterfactuals. What cases like JOSEPH show is that responding to this challenge by artificially changing the portion of the past that is held fixed is not going to help. Given the setup, holding the past fixed in JOSEPH in order to get rid of supposedly irrelevant differences between worlds (OPTION TWO), radically changes COMPARE's output on the value of death.

Adherents of COMPARE sometimes allow flexibility with respect to what actual events should be held fixed before considering counterfactuals.¹⁰ But no matter how flexible we are in this respect, COMPARE fails to concurrently accommodate both the fact that, in a normal context, Joseph's death is considered extremely harmful for him,

⁷ For generating the argument, it suffices that there is a context in which Joseph's death is extremely harmful (roughly to the degree designated by COMPARE and OPTION ONE minus the degree to which the pain involved in dying was harmful) and that COMPARE fails to establish its degree of harm.

⁸ I will not consider theories that give ultimate priority to factual similarity (including in facts obtaining after the antecedent event) in determining nearness relations. Not only because they are independently problematic, but because when coupled with such a theory, COMPARE obviously fails as an account of harm. Put crudely, the problem is that given the similarity of subsequent events, any event turns out hardly consequential for the subject on this view.

⁹ Theories of counterfactuals vary with respect to how they balance factual similarity and nomological similarity. Some theories give ultimate priority to nomological similarity, keeping the laws of nature fixed across possible worlds (under consideration). (This view has recently been defended by Dorr (2016). In the context of accounts of harm, it is held by McMahan (1988, p. 48). Other theories recognize counterfactual miracles: allowing nearby worlds to contain local, preferably small, divergences from the laws of nature so as to prevent E from occurring shortly before it does. This is Lewis's (1973) view.

¹⁰ Explicitly in the context of backtracking: McMahan (2002, 112–117) supposes that an intuitive comparison is merely a matter of correct tradeoff between similarity criteria. The analysis of cases like JOSEPH undermines this view.

and the fact that the pain he suffers before he dies should not contribute to the degree of harm assigned to his death.¹¹

3 Discussion

Before considering the proper reaction to the problem, let us examine some possible objections on behalf of the counterfactual comparative view.

First: One may worry that OPTION ONE and OPTION TWO are not exclusive candidates for describing the nearest possible world in which Joseph does not die the way he does.¹² In response: it should be noted that many alternative options will not make any important difference. For instance, if one judges that the nearest possible world in which Joseph does not die is such that Joseph *is* picked by the lottery but someone else is poisoned instead, COMPARE will still misleadingly imply that the harm of Joseph's death includes the suffering caused by the poison. Furthermore, if some possibility in which Joseph fails to die is (1) closer to actuality than OPTION ONE and OPTION TWO and (2) makes a difference with respect to our judgment, we will likely be able to make that possibility very remote by slightly modifying the case. I will not try to foresee every possible objection along those lines, but here is an example: if one finds the possibility that the commander will not carry out his threat even if Joseph miraculously fails to die worth considering, we can make this possibility remote by slightly modifying JOSEPH e.g. by adding that as a matter of policy, the army makes sure that its commanders always carry out their threats.

Second: One may argue that while COMPARE does not establish the degree to which Joseph's death is bad, it does establish the degree to which the compound event of his being poisoned and dying as a result is bad, and that this compound event is what we ought to care about when evaluating actual cases of death. I am sympathetic to the view that in many practical contexts 'dying' (i.e. death and the process leading to it) is categorically more basic than 'death', and is the more natural candidate for comparative evaluation (Pitcovski 2019). Nevertheless, COMPARE was introduced as a general account of the overall harm of particular events. A successful account ought to have left us the option to determine the value of Joseph's death if we were practically interested in it. It is true, perhaps, that sometimes the smallest event that can be said to give rise to some harm is a compound event and that in those cases COMPARE should not be expected to establish the harm of each component. But this is not the case when some component has a degree of harm that is clearly distinguishable from the harm of the compound event. So, for example, when the compound event is partly bad because it gives rise to intrinsically bad

¹¹ For this reason, emphasizing that assigning degrees of harm is context dependent along the lines of Bradley (2009), or Bradley (2004, pp. 49–50) will not help. If there is any context in which Joseph's death should appropriately be considered bad for him, COMPARE would fail to establish the degree to which it's bad for him in that context. Intuitively, for instance there is a perfectly normal context in which it is bad for Joseph to die at a young age. But the nearest possible world in which Joseph dies much older is either a world in which he fails to suffer from the poison, or a world in which he is tortured with the rest of his town.

¹² Given the dialectic, candidates based on an approach that renders COMPARE unappealing altogether are set aside. For this reason I do not consider the option on which Joseph dies of the poison a few milliseconds later than he actually does (see footnote 4).

states, while the harm of the component in question is strictly deprivational, COMPARE should be expected to establish the degree to which this component is harmful.¹³ This is exactly the case in JOSEPH: the value of Joseph's death remains the same independently of whether his being poisoned is pleasurable or painful (during the time preceding his death), while the value of the relevant compound event of dying is sensitive to that. An account that correctly establishes the value of Joseph's death is still called for.

Lastly: Adherents of the counterfactual comparative account might try to respond to the problem of backtracking (i.e. the problem created by the fact that relevant possible worlds have a different past, with independent value for the subject) by proposing a different comparative strategy. Most straightforwardly, they can associate the harm of an event with the comparative badness of what happens after it occurs. In a nutshell, they can replace COMPARE with:

COMPARE+: The degree of harm for S of an event E occurring at some time t is the degree to which the time from t onwards is worse for S in the actual world than in the nearest possible world in which E fails to take place.¹⁴

Unfortunately, this comparative strategy fails in cases where the nonoccurrence of events preceding death (in the relevant nearby world) has long-term consequences. Consider, for example:

INA: Ina was a distinguished, very young medicine student. She almost graduated to become the youngest MD ever. But then, as a result of drinking from a beer-glass she mistakenly took to be her own, Ina was infected with X-disease. The disease caused her terrible damage, preventing her from taking the exams and from becoming the youngest MD, and eventually caused her death.

Suppose that if Ina had not been infected (at t_0), she would have become the youngest MD ever (at t_1). Suppose that this fact alone would guarantee special benefits much later in Ina's life (at t_3). Presumably, in the nearest possible world in which Ina does not die at t_2 (between the time of the exams and the time of the expected benefits) Ina is not infected with X-disease, and she passes the exams at t_1 . If so, COMPARE+ fails to capture the degree to which Ina's death is harmful: That world is too good (for Ina) at times subsequent to t_2 . According to COMPARE+ Ina's death

¹³ The case of backtracking is essentially different from cases of overdetermination (and preemption) in this respect. Standard approaches to overdetermination are inapplicable to cases of backtracking. For instance, according to Feit (2015): E harms S to degree n iff E is the smallest super-plurality of every plurality of events P such that (1) if none of the events in P had occurred, S would have been better off by n, and (2) there is no smaller sub-plurality of P such that if none of the events in it had occurred, S would have been better off by n. In the case of overdetermination it is natural to take a plurality of events to jointly constitute a compound harmful event, E, since no individual event is such that if it wasn't for its occurrence, S would have been better off. However, in cases like JOSEPH the harm of some component, viz. Joseph's death, is intuitively different to that of the compound event. Formulated in terms of a counterfactual comparative account Feit's criterion fails exactly where COMPARE fails.

¹⁴ Aside for the problems discussed below, COMPARE+ seems to rely on the controversial presupposition that persons that no longer exist have degrees of well-being (particularly when this principle is applied to evaluate events of death). I will set the discussion of this potential problem aside.

is harmful partly for depriving her of the benefits due to graduating as the youngest MD. But the fact that she did not graduate as the youngest MD, and so was not expected to receive the benefits of being the youngest MD, has nothing to do with her death. It would be false to say that death deprived Ina of something she was not expected to receive given how her life unfolded before she died. Hence, it is false to say that death deprived Ina of the special benefits in question, or that it was harmful in this respect.

One might be tempted to argue that the nearest world in which Ina does not die is a world in which she is somehow cured at the time she was supposed to die. But there are two problems with this suggestion. First, we can always construct the case such that the miracle of being cured has to be enormous and to require heavy violations of fundamental laws of nature, while being infected is an extreme case of bad luck, so that things could easily have been different in this respect. Second, we can combine insights about JOSEPH and INA to construct a case in which any view about counterfactuals on which the subject is cured right before their death gives obviously wrong results when coupled with COMPARE+. Consider:

JOSEPHINA: Josephina was a distinguished, very young medicine student. She almost graduated to become the youngest MD ever. But then, as a result of drinking from a beer-glass she mistakenly took to be her own, she was infected with X-disease. The disease caused her terrible damage, preventing her from taking the exams (3 days later) and becoming the youngest MD. Given the laws of nature, no one can survive X-disease for more than 10 days. Josephina knew this. Near the end of day 10 of her disease, her town is captured by a cruel and powerful army. The people of the town are given the following choice: either they present a young fresh body outside the walls by next morning, or the whole town will be sentenced to a lifetime of torture. Although there are other young people in her situation, Josephina volunteers to be put in a cage outside the walls. She dies during the night, as expected.

As the case is constructed, if we suppose that the nearest world in which Josephina does not die is such that she is miraculously cured, according to COMPARE+ her actual death is not harmful, but beneficial; it saves her from a lifetime of torture. In some normal contexts, this seems to be the wrong result. COMPARE+ fails to establish that at least in some context, Ina's death is terrible for her while the degree to which it is bad has nothing to do with missing out on the special benefits of being the youngest MD.

4 Harm explained

A successful account of harm has to distinguish the harm that ought to be associated with death from harms having to do with the suffering preceding death or harms that are otherwise clearly unrelated to death. Before exploring the beginnings of an alternative, it would be helpful to have a clearer view of why COMPARE and CAUSE fail.

Note first that the problem we have encountered with COMPARE is that it counts *too many* states as constitutive of the value of an event (or its degree of harm). In the case of JOSEPH, not all of the states that fail to take place in the nearest possible world in which Joseph doesn't die seemed to be relevant for evaluating the harm of Joseph's death. While CAUSE is often regarded as too restrictive, COMPARE ought to be regarded as too permissive. CAUSE fails to account for deprivational harm, thereby implying that no case of death is harmful for its subject. COMPARE, on the other hand, considers every state that is counterfactually dependent on an event to contribute to its value, thus implying that the harm of some cases of death is not strictly deprivational.

CAUSE and COMPARE fail to capture the extension of our notion of harm. Some progress can be made if we can find common ground relative to which CAUSE is too restrictive and COMPARE is not restrictive enough. A promising candidate for providing this common ground, I suggest, is the notion of explanation. Crudely speaking, it may be argued that CAUSE is too restrictive because what is explained by the occurrence of death is not exhausted by what death causes; it also includes what death prevents, and what it constitutes. COMPARE, on the other hand, is too permissive because (put crudely) it sometimes counts states that are not explained by the occurrence of some event E to be part of the harm of E. While the occurrence of Joseph's death explains why he missed out on some pleasurable moments, for instance, it does not explain why he suffers terribly from the poison that kills him. COMPARE fails to make the distinction, allowing both what is explained by the occurrence of Joseph's death and what is otherwise counterfactually dependent on it to contribute to the value of Joseph's death.

The proposed analysis for the failure of COMPARE can further be tested by examining the outputs of COMPARE in other cases in which counterfactual dependence departs from explanation. Consider some particular event—the construction of a wall, for instance—that harms a person, Noah, by blocking the view from his apartment. Now suppose that the nearest possible world in which this event (the construction of the wall) fails to take place is also the nearest possible world in which some other events fail to take place (they may include: the appearance of the wall's shadow, the blossoming of a tree that needed that shadow etc.). Given counterfactual dependence relations, COMPARE implies that these further events (the blossoming of the tree etc.) harm Noah to the extent that the construction of the wall harms him. But this is clearly the wrong result.¹⁵ On the current analysis the source of the problem is that while the harm in question is counterfactually dependent both on the construction of the wall and on these other events, only the construction of the wall explains it.

In light of the above, it seems that in order to overcome the problems discussed so far, we need to base an account of harm roughly along the following lines:

EXPLAIN: An event E is a harm for S iff the totality of states that obtain (/ fail to obtain) *because* E occurs are overall intrinsically bad (/good) for S. The

¹⁵ For a similar objection to COMPARE see Purves (2019, penultimate paragraph of Sect. 3). Horwich (1989, pp. 169–170) and Maudlin (1994, pp. 128–129) discuss cases in which counterfactual dependence departs from causation, including cases of backtracking, and effects of a common cause. (Thanks to Arnon Keren).

degree to which an event E is harmful for S is the degree to which the states that obtain (/fail to obtain) *because* E occurs are overall intrinsically bad (/good) for S.

As opposed to COMPARE, EXPLAIN gives us the correct result about the harm of death in cases like JOSEPH and INA. Irrelevant states of suffering preceding death are not considered part of the harm of death, for while being counterfactually dependent on death, these states are not explained by it. Likewise for states that are prevented by something other than death. Moreover, unlike COMPARE, EXPLAIN does not imply that the wall's shadow is harmful to Noah (in the case discussed above): while some bad states are counterfactually dependent on this shadow, presumably, none is explained by it.

For all that has been said, EXPLAIN might not be the only account to meet these desiderata. Among accounts in the literature, EXPLAIN is in important respects most similar to a refined causation based account, according to which an event E is harmful for a person S to the extent that everything that E causes is intrinsically worse for S than everything that E causes not to happen (see Conee 2006, pp. 183–185), i.e.:

CAUSE+: An event E is a harm for S iff the totality of states that E causes to obtain (/fail to obtain) are overall intrinsically bad (/good) for S. The degree to which an event E is harmful for S is the degree to which the states that E causes to obtain (/fail to obtain) are overall intrinsically bad (/good) for S.

CAUSE+ nicely handles deprivational harms and the problem discussed in Sect. 2 (i.e. cases like JOSEPH): Unlike CAUSE, CAUSE+ establishes that events that cause good states failure to obtain give rise to deprivational harms. Unlike COMPARE, CAUSE+ does not misleadingly regard the pain preceding death as part of the harm of death in cases like JOSEPH, since that pain is not caused by death. However, EXPLAIN has other advantages. The main difference between CAUSE+ and EXPLAIN has to do with the fact that as opposed to CAUSE+ (or CAUSE, for that matter), EXPLAIN is not restricted to causal explanations. EXPLAIN is in principle open to the idea that events can generate harms by standing to states in explanatory relations other than causation and prevention. This makes it possible for EXPLAIN to account for harms that are constituted by an event rather than being caused by it, for example. Even if my joining the cult constitutes my mother's desire frustration, rather than causing it, my joining the cult explains my mother's desire frustration. It therefore harms her by the lights of EXPLAIN, but not by the lights of CAUSE+.

Recapitulating, EXPLAIN seems to be on the right track. (1) Unlike CAUSE, EXPLAIN accounts for deprivational harms, including the harm of death: it takes the value of any state that fails to obtain because of death into account when the value of death is considered. Moreover, since it is not restricted in advance to causal explanations, unlike CAUSE or CAUSE+, EXPLAIN also nicely accounts for harms that are constituted by an event rather than caused by it. (2) As opposed to COMPARE, EXPLAIN gives us the correct result about the harm of death in

cases like JOSEPH and INA. Irrelevant states of suffering preceding death are not considered part of the harm of death, for while being counterfactually dependent on death, these states are not explained by it. Likewise for states that are prevented by something other than death. Moreover, unlike COMPARE, EXPLAIN does not imply that the wall's shadow is harmful to Noah (in the case discussed above) for while some bad states are counterfactually dependent on this shadow, presumably, none is explained by it.

EXPLAIN still needs to be unpacked and further examined. The main notion that needs to be unpacked is 'because'. As I use this notion it should be understood strictly as an explanatory locution. This accords with standard philosophical usage. 'Because' is a sentential connective, binding together *explanandum* and *explanans* sentences. Unlike 'caused' it does not designate a relation holding between particulars, e.g. particular events. Since EXPLAIN deals with the harm associated with events, the *explanans* in this case takes a particular form, 'E occurs'. Since we have been assuming all along that the vehicles of intrinsic goodness or badness are states,¹⁶ we will be looking at every *explanandum* sentence taking the form 'such-and-such a state obtains' that forms a true sentence when it fills in the blank in the formula '____ because E occurs' (where E is the event we are trying to assess). To avoid ungainly formulations, I will allow myself to use phrases of the form 'every state that obtains *because of E*' as shorthand for 'every state X of which it is true to say 'X obtains because E occurs''.

Two further clarifications are called for: (1) It should be noticed that whether the occurrence of E explains the obtaining of some state may depend on the description of E. But whether E harms an agent does not depend on how E is described. In other words: while 'because' is plausibly hyperintensional, 'harms' is definitely not. To avoid this mismatch, we will give every description a *de re* reading, so that it picks out a particular event with a particular modal profile. When considering what is explained by 'E occurs' we look at what can be explained by all sentences that replace 'E' with a description that picks out an event with the same modal profile. (2) Some philosophers distinguish between 'explanation' and '*adequate* explanation'. For instance, those who take an event to explain much of what is counterfactually dependent on it, will want to distinguish between 'Beth's flowers died because Beth failed to water them' and 'Beth's flowers died because the Queen of England failed to water them'. In a context in which the latter sounds absurd, they will want to say that while there is some sense in which the Queen's failure to water the plants explains their death, given background presuppositions (perhaps about the difference in responsibilities, and about what is normal), only Beth's failure to water the plants *adequately* explains their death.¹⁷ In case there is such a distinction, in what follows I will use 'explains' to mean '*adequately* explains' (without commitment to any particular account of adequate explanation). I will emphasize this when required.

¹⁶ Although nothing said so far makes this presupposition committal. EXPLAIN can be modified to an account that takes events as vehicles of intrinsic harm (in the spirit of Hanser (2008), for instance), or an account that takes facts to play this role.

¹⁷ See Beebe (2004, pp. 307–8) for instance.

Extensions of EXPLAIN will reasonably be called for. Given the last part of the previous paragraph, one obvious requirement is the need to account for cases of harm by omission, i.e. cases in which a possible event's failure to take place is a source of harm. Explanation based accounts can easily account for such cases. For example: If Allan's neighbor Beth was supposed to water Allan's plants when he was away, her failure to do so is a source of harm to the extent that it explains why the flowers in Allan's garden have died (and to the extent that their death intrinsically harms Allan). This extension can be accommodated without metaphysical commitment to events of nonoccurrence. Since fewer unnecessary commitments better serves the main purpose, we can follow Collins (2000) for instance¹⁸ and, rather than speaking of events of nonoccurrence (like 'the event of not watering the plants'), simply speak of the nonoccurrence of events. We can match every possible event, E, to its occurrence proposition, the proposition expressed by 'E occurs', and to its nonoccurrence proposition, expressed by 'it is not the case that E occurs'. We can then allow that the right-hand operand (of 'because') reserved for the *explanans* to take as values sentences expressing either an occurrence proposition or a nonoccurrence proposition. In other words, given its specific interest, EXPLAIN is restricted to cases in which the occurrence of an event or the nonoccurrence of an event explains the obtaining of a state or its non-obtaining. (I will avoid the heavy-handed formulation of the extended version of EXPLAIN).

An attractive feature of EXPLAIN is that it can remain neutral with respect to accounts of explanation. As an account of harm, it merely requires minimal features of explanations. Roughly (1) that not everything counterfactually dependent on an event is thereby explained by it (no commitment to any specific account of explanatory relevance is made). (2) That the 'because' relation underlies relations integral to overall harm, for instance, that when an event brings about or constitutes (or prevents) some state, it thereby explains it (or its absence, respectively).

5 EXPLAIN defended

EXPLAIN nicely clarifies why our notion of harm plays a key role in the evaluation of events we may be subjected to, and of alternative courses of actions we plan to undertake: If I, or someone else, would suffer (all in all) bad states *because of* some course of action I am deliberating about, this is a good *pro tanto* prudential reason, or moral reason (respectively) to avoid it. Moreover, EXPLAIN commands general theoretical virtues, like simplicity: it provides a unified account of harm, rather than an ad hoc list of conditions; it is theoretically and metaphysically parsimonious, as it does not rely on a specific metaphysical view or account of explanation. EXPLAIN also meets other, more specific requirements of accounts of overall harm¹⁹: It remains amoral and neutral in the sense of respecting the distinction between harming and being morally blameworthy and in allowing events other than agential actions to count as harms (accordingly). It does not presuppose any theory of well-being, as required

¹⁸ Beebe (2004) defends the view that absences (like the failure of an event to take place) do not cause, while respecting the role absences play in our explanatory practices.

¹⁹ See Bradley (2012), (Sect. 2), Purves (2019, p. 2645).

of an account of *overall* harm. It supports a plausible account of agential harm when restricted to agential actions, and it allows for a symmetrical treatment of harm and benefit. If harms and benefits require symmetric accounts, a question that will not be decided here, an extended version of EXPLAIN can include the following: “An event E benefits S to the extent that overall, the total states that obtain (/fail to obtain) because E occurs are intrinsically good for S (/bad for S, respectively, in the case of events that fail to obtain)”.

In the remaining paragraphs I will briefly demonstrate how, on top of meeting minimal requirements of an adequate account of harm, EXPLAIN can meet our intuitions about issues that often figure as a challenge for theories of harm: (1) cases of redundant causation (for conciseness, I focus on cases of preemption), (2) the interplay between context and judgments of harm, (3) the distinction between harming and failing to benefit, and (4–5) two worries originally raised against CAUSE+. My aim in those paragraphs is not to provide a full explication and defense of EXPLAIN, that would require far more space than I have available, but rather: to consider it as an alternative to existing accounts, and to argue that it is at the very least worthy of further consideration.

(1) Cases of preemption are cases in which some event brings about an outcome, where another event (which does not actually bring it about) would have brought it about, in case the former event hadn't. This is what happens when the shot of one sniper kills me first, but another shot of a different sniper would have killed me if the first shot hadn't. Such cases challenge views like COMPARE: since it seems that my death is not counterfactually dependent on any one of the shots, as I am not much better off in the nearest possible world in which any one of the shots fails to take place, how can we say of any shot, in particular: of the shot that actually killed me, that it gravely harmed me on such a view?

EXPLAIN seems to provide an easy response: While the deprivational harm associated with my death is not counterfactually dependent on the first sniper's shot, it is nonetheless explained by it. The shot that hit me when I was dead, on the other hand, fails to explain my death, or the deprivational harm associated with it.

Since not all theories of explanation analyze cases of preemption in this way,²⁰ and since I wish to remain neutral with respect to the correct account of explanation, it is important to emphasize that the means developed by adherents of the counterfactual comparative view are also available upon slight adjustments to adherents of EXPLAIN. Hence, one possible response is to implement Feit's (2015) notion of plural harm, originally introduced for expanding COMPARE in response to this challenge. Consider:

PLURAL-EXPLAIN: E harms S to degree n iff E is the smallest super-plurality²¹ of every plurality of events P such that (1) The total intrinsic value of

²⁰ Adherents of the network model of causal explanation may opt for PLURAL-EXPLAIN discussed below. This fits nicely with the way it handles redundant causation. See Paul and Hall (2013, ch.3) for an elaborate discussion.

²¹ A plurality of events E1, E2 is nothing over and above those two events, and a super-plurality of events is nothing over and above the events which some pluralities, its sub-pluralities, consist of. (So a super-plurality of the plurality that includes E1 and E2, and the plurality that includes E2 and E3, will include precisely E1, E2 and E3). This technical notion is used in order to cover cases where (for example) every two events (among E1, E2 and E3) would suffice to bring about a harm of degree n, but where all three events take place.

states explained by E for S is n; and (2) there is no smaller sub-plurality of P such that the intrinsic value of states explained by it is n.²²

Suppose that the state of being dead is not explained strictly by the first shot that hit me. (The first shot on its own fails to explain anything that bad. Given the second shot, the first shot only explains why I missed out on a few moments of life). In this case PLURAL-EXPLAIN implies that no single shot harms me, while also implying that both shots jointly harm me to an enormous degree because they jointly explain why I failed to go on with my life. (This accords with Feit's analysis of cases of preemption and overdetermination).²³

Since no adequate theory of explanation would allow that my death is not *at least partly* explained by the shot that killed me, an explanation based theory of harm can either respond to the problem of preemption by straightforwardly accounting for the harm associated with the first shot, or by adopting PLURAL-EXPLAIN. In any event, such views will fare no worse than COMPARE in this respect.

(2) A second issue is that judgments of harm vary with conversational context. For instance, when looking at JOSEPH, we normally think of Joseph's death as terrible for him: Joseph failed to go on with his life projects because he died. But there may be conversational contexts in which we would not think of Joseph's death as very bad for him. For instance: a context in which Joseph's poisoning is held fixed. For example when he is already outside the walls, and two of his townmates are discussing whether it is good for him to die. However, while changing the context changes our judgment of harm, it correspondingly changes what is explained by death. In the latter context, Joseph's death no longer explains why he fails to go on with his life projects. It only explains why he avoids a lifetime of torture.

Likewise, there may be conversational contexts in which we would be inclined to associate the pain that takes place before death with the harm of death.²⁴ This is what happens when we specify Joseph's death by comparing it to another possible death, which is similar but much quicker, thereby involving less pain, as it were. We may be inclined to associate the enduring pain with the fact that Joseph suffered the actual delayed death (rather than the possible quicker death). And indeed, there may be a context in which the assertion 'Joseph's delayed death is bad in part because it involves more suffering' expresses a truth, rather than merely communicating 'Joseph's enduring *life* was miserable partly due to living long after taking the poison'. However, in such contexts, we would also be inclined to say that it is *because* his death was delayed that Joseph suffered (i.e. that a feature of his death, namely its being delayed, explains why he suffered for longer).

In order for context-dependent judgments to generate a challenge for EXPLAIN, there needs to be a single context in which there is a mismatch between the aggregated intrinsic value of states (and non-obtaining states) that are explained by E,

²² Feit's original criterion is discussed in footnote 13.

²³ Carlson et. al. (Manuscript) challenge accounts that address problems of redundant causation by covering harm of pluralities of events. PLURAL-EXPLAIN, like Feit's (2015) view, will have to face the challenges they raise.

²⁴ Bradley (2004, p. 11) calls such contexts 'non-standard'.

and the overall harm of E. Yet the cases of context-shift discussed in the literature on harm are cases in which context-shifts change our judgment about the harm we associate with an event by correspondingly changing what is explained by the event in question.²⁵

(3) Some theorists of harm distinguish harm from failure to benefit. This is thought to create a problem for accounts like COMPARE. Consider:

BATMAN: Suppose that Batman purchases golf clubs with the intention of giving them to Robin, but the Joker persuades him to keep them for himself. Had Batman not kept the clubs he would have given them to Robin. (Bradley, 2012, p. 397)

and.

JOKER: Batman has delivered golf clubs to Robin, and the Joker removes the clubs just before Robin opens the door to retrieve them. Had the Joker not removed the clubs from Robin's porch, Robin would have found them and would thereby be better off. (Purves, 2019, p. 2633)

COMPARE implies that the Joker's actions in JOKER and Batman's behavior in BATMAN both harm Robin. Yet intuitively, only the Joker's actions in JOKER actually harm Robin. Batman's failure to give the clubs to Robin in BATMAN is not a harm but, at best, a failure to benefit.

Responding on behalf of EXPLAIN seems easy: as noted earlier (while discussing the case in which the queen fails to water Beth's plants), part of what is required for an event, or an action, to adequately explain some outcome depends on normal presuppositions, background obligations etc. While it is clearly a problem for COMPARE if BATMAN is not a case of harm, it is only a problem for EXPLAIN if Batman's behavior adequately explains the non-obtaining of some states that are intrinsically good for Robin (in this case). This is far from obvious.

In any event, a counterexample to EXPLAIN should hardly be expected by cases of this sort, for while intuitions about whether BATMAN is a case of genuine harm do pull in different directions, intuitions about whether it is a case of adequate explanation change accordingly. Hence, if by some account of explanation Batman's behavior adequately explains Robin's failure to benefit in BATMAN, adherents of this account ought to be disposed to follow Hanna (2015) in judging that Batman does harm Robin in BATMAN.²⁶ The view that Batman harms Robin but that the states constituting this harm are not adequately explained by Batman's behavior seems highly implausible. The view that states that obtain because of Batman's behavior are altogether intrinsically harmful to Robin, but that Batman fails to harm Robin, is equally implausible.

²⁵ This includes cases in which context picks out a different *explanans* by changing the way in which events are individuated. (see McMahan (1988, p. 45) for instance).

²⁶ In response to this challenge Purves upholds Woollard's (2008) distinction between making and allowing to argue that only what we make (not what we allow) constitutes the degree of harm we bring about (See Purves (2019, p. 2644) for a concise presentation of the distinction). See Johansson and Risberg's (2020) criticism.

(4) Two challenges originally leveled against CAUSE+ may seem to arise for EXPLAIN as well. The first challenge is to specify which things exactly are *caused* not to exist by an event, or, in the case of EXPLAIN, to specify the states that fail to obtain *because* of an event. This is important if we are to determine degrees of deprivational harms. More specifically, the worry is that one event can seemingly explain the non-obtaining of indefinitely many incompatible states (Feit 2015, pp. 365–6). For example, S’s death in a car accident at 14:00, just before entering an amusement park, can explain the non-obtaining of indefinitely many states of S with varying degrees of pleasure at exactly 15:00, supposing that the premature death explains why S was not on any of the attractions at 15:00 (and provided that different attractions give different degrees of pleasure). EXPLAIN seems committed to taking the value of all incompatible states that fail to obtain because of E to contribute to the degree of harm associated with E. This is a problem.

One way to respond to this worry is to adopt an account of explanation that can coherently make sense of accepting that E explains why it is not the case that either State₁ obtains or State₂ obtains or... (henceforth ‘E explains \sim (State₁ \vee State₂...)’) while denying that E explains why State₁ fails to obtain and E explains why State₂ fails to obtain... (i.e. denying that E explains \sim State₁ & E explains \sim State₂). A promising way of doing so is to adhere to a contrastive account of explanation, such that in some cases E explains \sim (State₁ \vee State₂...) rather than (State₁ \vee State₂...) without it being the case that E explains \sim State₁ rather than State₁, for instance. (see Botterill (2010), van Fraassen (1980). In the example above, the premature death does seem to explain why S was not on any of the attractions rather than being on one of them. However, given the indeterminacy of specific whereabouts that initially gave rise to the puzzle, S’s death falls short of explaining why S was not, for example, on the hot air balloon at 15:00 rather than being on it, S’s being on the hot air balloon at that time requires more than not being dead in the setup. Given the connection between particular attractions and particular states, letting State_{a1} stand for the pleasure of being on the hot air balloon, S’s death explains \sim (State_{a1} \vee State_{a2}...) rather than (State_{a1} \vee State_{a2}...) without it being the case that this death explains \sim State_{a1} rather than State_{a1}.

Other options are open for EXPLAIN theorists. In principle, EXPLAIN is hospitable to accounts of explanation that take counterfactual dependence to be essential for explanation. In addressing the case discussed above, such accounts can say that the premature death explains the non-obtaining of a certain state only if it obtains in the nearest possible world in which this case of death doesn’t take place. In case the state in question obtains only in some of the relevant nearest possible worlds, we get indeterminate value (which is exactly what COMPARE establishes successfully, in my view, in the case above (see Pitcovski and Peet 2022)). Any account along those lines should acknowledge (following Bradley 2012, p. 409) that if the states that fail to obtain because of E are just the ones that would have obtained if it weren’t for E, then we are back to a counterfactual account of harm. Hence, such views are required to distinguish ‘explained by E’ from ‘being counterfactually dependent on E’ in a non ad hoc way that preserves the advantages of EXPLAIN with respect to the problem of backtracking. This can be done by relying on the fact that laws underlie causal and constitutive relations but not counterfactual dependence relations in general (see

Maudlin 2007, ch.5). Hence, an event explains only (but not all of) what is counterfactually dependent on it. One shortcoming of this strategy is that while providing a satisfactory response to the problem of backtracking and to the amusement park case, it gives up EXPLAIN's advantages over COMPARE with respect to the other problems discussed in this section (flexibility with respect to the problem of preemption, for instance). This is not a decisive reason for rejecting it, but it is definitely a price to pay. While contrastivism seems more neutral and promising at this stage, exploring the pros and cons of each strategy in full detail will have to wait for another time.

(5) An independent challenge for EXPLAIN, which was also originally leveled against CAUSE+, is based on special cases of preemption that seemingly give rise to judgments of overall harm that conflict with EXPLAIN. Consider the following case by Carlson et al., (2021, p. 8):

RIDDLER: Riddler is about to spray tear gas in Batman's left eye. The Joker can prevent this, either by simply telling Riddler to leave Batman alone, or by spraying tear gas in Batman's right eye. (Riddler and the Joker have agreed to leave at least one of Batman's eyes undamaged.) If the Joker tells Riddler to leave Batman alone, no tear gas will be sprayed. If Batman gets tear gas in his left eye he will suffer 15 units of pain. If he gets tear gas in his right eye he will suffer 10 units of pain.

According to Carlson et.al., "Intuitively, the Joker's telling Riddler to leave Batman alone would benefit Batman overall, while spraying would harm him overall." (ibid). Nevertheless CAUSE+, and indeed EXPLAIN as well, "implies that [the Joker's] action would benefit Batman overall. Although it would cause Batman's suffering 10 units of pain to obtain, and hence be *pro tanto* harmful, this harm is outweighed by the benefit constituted by Batman's not suffering 15 units of pain."

One way to respond to this worry is to straightforwardly reject the intuition that spraying Batman, i.e. preventing 15 units of pain while causing 10 units of pain, overall harms Batman. This response is always open for adherents of EXPLAIN and it is not altogether untenable or absurd, given that intrinsic harm is outweighed by intrinsic benefit in this case. Nevertheless, some versions of EXPLAIN are able to explain the sense in which the intuitions underlying the original judgment are also acceptable. This can be done by expanding the contrastive account of explanation so as to allow contrastive *explanantia* (perhaps, but not necessarily, in addition to traditionally allowing contrastive *explananda*).²⁷ This way, EXPLAIN theorists will be able to confirm 'The Joker's spraying Batman rather than telling Riddler to leave him alone overall harms Batman' while denying that 'the Joker's spraying Batman rather than not spraying Batman harms Batman'. The first contrastive *explanans* explains why Batman suffers pain at all, and this is the sense in which the Joker's choice is overall harmful for Batman. The second contrastive *explanans* explains why Batman fails to suffer 15 units of pain (despite also explaining why he suffers 10 units) and this is a sense in which the action is overall beneficial for Batman. So

²⁷ Schaffer (2005) takes causation to be contrastive in this twofold way.

the precise sense in which the intuition about the Joker's action being an overall harm for Batman is correct can be comprehensively established.²⁸

Aside for addressing the problem discussed in Sect. 2, EXPLAIN does at least as well as its alternatives regarding our intuitions about the extension of harm. This gives further support to EXPLAIN.

6 Conclusion

By the eyes of an explanation based approach, causation based accounts and COMPARE both seem to capture something important about harm.²⁹ Since much of what is explained by an event is also caused by it, the intuitive appeal of causation based accounts is unsurprising. But as noted, causation based accounts fail to classify harms having to do with deprivation or constitution as harms. Taking this into account, COMPARE was designed to cover all cases of harm without going ad hoc. However, as we have seen, COMPARE sometimes erroneously takes the harm of an event to consist of states that are not in any appropriate way related to it. Since EXPLAIN guarantees a link of relevance between an event and the states comprising its harm, and since it also seems to cover the dead area of deprivational and constitutional harms, EXPLAIN ought to be further considered.³⁰

References

- Beebe, H. (2004). Causing and nothingness. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 291–308). MIT Press.
- Botterill, G. (2010). Two kinds of causal explanation. *Theoria*, 76, 287–313.
- Bradley, B. (2004). When is death bad for the one who dies? *Noûs*, 38(1), 1–28.
- Bradley, B. (2009). *Well-being and death*. Oxford University Press.
- Bradley, B. (2012). Doing away with harm. *Philosophy and Phenomenological Research*, 85(2), 390–412.
- Broome, J. (1999). *Ethics out of economics*. Cambridge University Press.
- Broome, J. (2004). *Weighing lives*. Oxford University Press.
- Carlson, E., Johansson, J., & Risberg, O., (2021). Causal accounts of harming. *Pacific Philosophical Quarterly*, 1–26.
- Carlson, E., Johansson, J., & Risberg, O., (2022). Plural harm: Plural problems.
- Collins, J. (2000). Preemptive Prevention. *The Journal of Philosophy*, 97(4), 223–234.
- Conee, E. (2006). Dispositions toward counterfactuals in ethics. In J.R. Raibley, M.J. Zimmerman, & K. McDaniel (Eds.), *The Good, the Right, Life and Death: Essays in Honor of Fred Feldman*, Ashgate, (p.173–187).
- Dorr, C. (2016). Against counterfactual miracles. *The Philosophical Review*, 125(2), 241–286.

²⁸ Likewise for the case discussed by Norcross (Norcross 2005, pp. 165–6) where intuitively there is a sense in which Bobby Knight harms the person he chokes, even if had he not choked her, he would have torn off both her arms. Contrastivists can say that Knight's choking the person rather than restraining himself explains why that person suffered pain. Nonetheless, Knight's choking the person rather than tearing off her arms explains why she is altogether well off.

²⁹ Given space limitations, other accounts were not addressed.

³⁰ Thanks to audiences at the University of Haifa and at the Hebrew University of Jerusalem for helpful discussions. Special thanks to Erik Carlson, David Enoch, Eric Olson, Andrew Peet, and anonymous referees of *Philosophical Studies* for their priceless comments.

- Feit, N. (2015). Plural harm. *Philosophy and Phenomenological Research*, 90(2), 361–388.
- Feldman, F. (1991). Some puzzles about the evil of death. *The Philosophical Review*, 100, 205–227.
- Hanna, N. (2015). Harm: omission, preemption, freedom. *Philosophy and Phenomenological Research*, 91(2), 1–23.
- Hanser, M. (2008). The metaphysics of harm. *Philosophy and Phenomenological Research*, 77(2), 421–450.
- Horwich, P. (1989). *Asymmetries in time*. MIT Press.
- Johansson, J., & Risberg, O. (2020). Harming and failing to benefit. *Philosophical Studies*, 177(6), 1539–1548.
- Klocksiem, J. (2012). A defense of the counterfactual comparative analysis of harm. *American Philosophical Quarterly*, 49(4), 285–300.
- Lewis, D. (1973). *Counterfactuals*. Blackwell.
- Maudlin, T. (1994). *Quantum Non-locality and relativity*. Blackwell.
- Maudlin, T. (2007). *The metaphysics within physics*. Oxford University Press.
- McMahan, J. (1988). Death and the value of life. *Ethics*, 99, 32–61.
- McMahan, J. (2002). *The ethics of killing: Problems at the margins of life*. Oxford University Press.
- Norcross, A. (2005). Harming in context. *Philosophical Studies*, 123, 149–173.
- Paul, L. A., & Hall, N. (2013). *Causation: A user's guide*. Oxford: Oxford University Press.
- Pitcovski, E. (2019). Jean Améry on the value of death and dying. In Y. Ataria, A., Kravitz, & E. Pitcovski (Eds.), *Jean Améry: Beyond the mind's limits* (pp. 221–236).
- Pitcovski, E., & Peet, A. (2022). Counterfactuals, indeterminacy, and value: A puzzle. *Synthese*, 200(1), 1–20.
- Purves, D. (2019). Harming as making worse off. *Philosophical Studies*, 176(10), 2629–2656.
- Schaffer, J. (2005). Contrastive causation. *The Philosophical Review*, 114(3), 327–358.
- Shiffrin, S. V. (2012). Harm and its moral significance. *Legal Theory*, 18(3), 357–398.
- van Fraassen, B. C. (1980). *The scientific image*. Clarendon Press.
- Woollard, F. (2008). Doing and allowing, threats and sequences. *Pacific Philosophical Quarterly*, 89(2), 261–277.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.