



Over-Determined Harms and Harmless Pluralities

Björn Petersson¹

Accepted: 25 July 2018 / Published online: 1 August 2018
© The Author(s) 2018

Abstract

A popular strategy for meeting over-determination and pre-emption challenges to the comparative counterfactual conception of harm is Derek Parfit's suggestion, more recently defended by Neil Feit, that a plurality of events harms A if and only if that plurality is the smallest plurality of events such that, if none of them had occurred, A would have been better off. This analysis of 'harm' rests on a simple but natural mistake about the relevant counterfactual comparison. Pluralities fulfilling these conditions make no difference to the worse for anyone in the over-determination cases that prompted the need for revising the comparative conception of harm to begin with. We may choose to call them harmful anyway, but then we must abandon the idea that making a difference to the worse for someone is essential to harming. I argue that we should hold on to the difference-making criterion and give up the plural harm principle. I offer an explanation of why Parfit's and Feit's plural harm approach seems attractive. Finally, I argue that the consequences of giving up the plural harm principle and holding on to the simple comparative counterfactual analysis of harm are less radical than we may think, in relation to questions about wrongness and responsibility.

Keywords Harm · Plural harm · Over-determination · Pre-emption · Accountability · Coresponsibility · Parfit · Feit

1 Harmless Pluralities

I just had this paper rejected by the *Journal of Over-Determination*. According to that journal's strict policy, manuscripts are rejected when one of the two reviewers suggests rejection, regardless of what the other reviewer says. In this case, both reviewers wanted the paper rejected.

✉ Björn Petersson
bjorn.petersson@fil.lu.se

¹ Department of Philosophy, Lund University, Box 192, 221 00 Lund, Sweden

The simple comparative counterfactual analysis of harm says that *an event E harms A if and only if A would have been better off had E not occurred*.¹

I would have been better off had this paper not been rejected. I seem to be harmed, and the two reviewers appear to bear some responsibility for that. However, assuming that the comforting effect of one positive review would have been outbalanced by the frustration created by being rejected in spite of such a review, I would not have been better off if just one of the reviewers had liked the paper. So, on the simple comparative counterfactual conception of harm, no single decision harms me.

There is a genre of less mundane stories illustrating over-determination or pre-emption² challenges to comparative conceptions of harming.³ These stories appear to undermine the common intuition that motivates the comparative counterfactual analysis: “that harms *make a difference*, in a negative way, to the person harmed”. (Bradley 2012, 397).

A popular way of meeting this challenge that feels “natural and extremely appealing” (Feit 2015, 369) is to assume that in cases like these we will have to regard the *set* (Parfit 1984, 70) or *plurality* (Feit 2015) of initial events as what harms. In Parfit’s terms, reviewers #1 and #2 *together* harm me. Parfit mostly puts the suggestion in terms of *groups* and *acts* (Parfit 1984, 71), but since we want a principle for dealing with harms that are over-determined not only by actions, but by any kind of events that can harm, it is preferable to formulate the principle in terms of events rather than acts.⁴ In Feit’s terminology:

A plurality of events harms A if and only if that plurality is the smallest plurality of events such that, if none of them had occurred, A would have been better off.⁵

I agree that this way of treating over-determination and pre-emption cases is intuitively appealing. However, a striking feature of this approach that has not been highlighted is that in the cases construed to challenge the simple comparative counterfactual account of harm, the plurality of events picked out by the analysis makes no difference to the worse for anyone.

What features are required for these cases to create problems for the simple counterfactual comparative account? In a real life case, the most obvious counterfactual assumption we would make is probably that if they had not both rejected, then either only

¹ Throughout this paper, I am talking about what harms a person *overall*. We could adopt the simple counterfactual comparative analysis of overall harm and still admit that there is a sense in which, say, my visit to the dentist is a pro tanto harm even though I would not have been better off overall if it had not occurred.

² In a pre-emption case, like in the over-determination case, the two events are each sufficient for the effect, but one triggers the effect before the other and pre-empts the causal role of the other event: X poisons me and Y then kills me before the poison would have killed me. (Parfit 1984, 70–71)

³ E.g.: X and Y shoot me simultaneously and each shot is lethal (Parfit 1984, 70), Batman has a lethal heart attack while being hit by a flaming cannonball (Bradley 2012, 406). See also Hanser 2008; Norcross 2005; Feit 2015. My reviewer example may seem less straightforward than some of these cases. Some might find it difficult to imagine that the reviews are submitted simultaneously or that the process is as rigid as claimed, etc. However, even Parfit’s seemingly cleaner example is under-described, and the reader is trusted with filling in the relevant details making it a challenge to the difference-making intuition about harm. I assume here that my reasoning about the plural harm principle as applied to the reviewer case applies to any of the standard examples of over-determined harms.

⁴ See Feit 2015, 370, and Petersson 2004, 305, about the importance of applicability to non-agential events.

⁵ Any act not affecting me, like some unrelated act by *Fred Astaire*, (Parfit 1984, 71) would, along with the acts of reviewer #1 and reviewer #2, belong to a set of events such that if none of them had occurred, I would not have been harmed. Hence the restriction to “the smallest group”. Feit refines the analysis further to cover more complex cases, but those elaborations do not affect any of the points I want to make here.

one of them would have rejected or none of them had done so. So, if they had not both rejected, the paper *might* have been accepted. If we apply the simple counterfactual comparative account in these circumstances, we would be allowed to say that R#1's act or R#2's act, or the plurality consisting of their two acts, *might* have harmed me. But this would merely reflect our epistemic position, and the case would not be very helpful when it comes to understanding under what conditions an event actually harms.

The cases under consideration are construed to help us examine conditions for *objective* harming. They are challenges to the simple counterfactual conditional account because that account fails to give a plausible verdict in a fully specified case where it is true by stipulation that (1) R#1 would have rejected the paper even if R#2 had accepted, and that (2) R#2 would have rejected the paper even if R#1 had accepted. These are the two conditionals that make it false to say that either R#1 or R#2 harms me on the simple counterfactual conditional account of harming.

From (1) and (2) it follows that if R#1 and R#2 had not both rejected the paper, one of them would have still done so. In Lewisian terms: (1) is true iff some world in which R#1 rejects but R#2 does not is closer to our world than any world in which neither rejects, and (2) is true iff some world in which R#2 rejects but R#1 does not is closer to our world than any world in which neither rejects. In other words, given the essential features of the initial case description, if the plurality of events consisting in the acts of R#1 and R#2 had not occurred, one of the elements in that plurality would still have occurred, and my paper would still have been rejected.

A plural harm account requiring negative difference-making would need a clause stating that a plurality of events harms A on the condition that A would have been better off had that plurality not occurred, and this clause would not be fulfilled in the cases under consideration. In the over-determined rejection case, the plurality of events consisting in the two reviewer recommendations is the smallest actual plurality of events such that, if none of them had occurred, I would have been better off. What makes the case problematic for the simple comparative counterfactual analysis of harm is the stipulation that if one of the reviewers had suggested acceptance, the other one had still rejected, which implies that if both had not rejected, one of them had still done so. This means that if the plurality had not occurred, sufficiently many of its elements would still have occurred, and the occurrence of the plurality therefore makes no difference to the worse for me.

If we think that the single events are harmless in the over-determination cases under consideration because they make no difference to the worse for anyone, we should then conclude that the pluralities are harmless for the same reason. If, on the other hand, we want to hold on to the view that a plurality of events harms A if and only if that plurality is the smallest plurality of events such that, if none of them had occurred, A would have been better off, we will have to give up the view that making a difference to the worse for someone is essential to harming. That would be to abandon the strong intuition that motivated the counterfactual comparative account, and it would mean that there is no challenge from over-determination and pre-emption to begin with.

The problems about harm discussed here should be separated from the problems that over-determination and pre-emption cases raise for certain theories of causation. Such cases are standard challenges to the claim that counterfactual dependence is necessary for causation, i.e. the idea that in order for something to be a cause of an event it must have made a difference to the occurrence of that event. However, even if a sufficiently subtle counterfactual account of causation, perhaps along the lines suggested by Lewis (2000), can deal with these challenges

in a satisfactory manner, this will not save the comparative counterfactual analysis of harm. In section 2, I briefly discuss and dismiss this possibility.

It may seem odd that those of us who share the intuition that negative difference-making is essential to harming tend to feel attracted by a plural harm principle that is incompatible with the difference-making criterion for harm. However, there are ambiguities in the question of what would happen if a plurality of events had not occurred, and it may be tempting to waver between interpretations when considering the analysis and the case-descriptions. This may explain why we tend to overlook the counterintuitive implications of the plural harm account. One source of confusion is that normally, when we ask what would have happened if two unrelated events had not both occurred, the issue of whether the relevant comparison is a world where one of them takes place or a world where neither occurs will vary with the context in which we ask the question. Another relevant ambiguity is between distributive and collective readings of expressions like “both events” and “a plurality of events”. Section 3 elaborates these diagnoses.

Section 4 suggests that there may be other reasons for finding acts or events bad and for holding people responsible on account of their acts than that they straightforwardly harm someone, even within a consequentialist framework.

2 A Note on over-Determination, the Counterfactual Comparative Notion of Harm, and the Counterfactual Dependence Condition for Causation

The counterfactual account of causation says that causing an event requires making a difference to the occurrence of that event,⁶ while the counterfactual comparative notion of harm says that harming someone requires making a (negative) difference to that person’s wellbeing. In the case under consideration, neither reviewer decision makes a difference to this article’s being rejected. If counterfactual dependence — making a difference to the occurrence of an event — is necessary for being part of the causal history of the event, it seems not only that none of the reviewers’ decisions harms me, but also that neither decision is part of the causal history of the rejection, which may be even harder to swallow. So, the problems affecting counterfactual dependence accounts of causation may seem analogous to those affecting the counterfactual comparative notion of harm.

The most promising strategy for meeting these challenges to the counterfactual dependence condition for causation is probably David Lewis’ move to allow for “fragile” versions of events, where a “version” of an event is a “tailor-made proposition about that event, that will be a proposition about how and when and whether the event occurs”. (Lewis 2000, 187) On this account, unlike in Lewis’ earlier counterfactual theory, it is not the *occurrence* of the caused event that is counterfactually dependent upon the causing event, but the exact *way* in which the event occurs. Although the coarsely individuated event ‘this article’s rejection’ would have taken place even if Reviewer #1 had liked the paper, it would not have occurred in exactly the same manner — the rejection letter would have been phrased slightly differently, my reactions would have been more complex, etc. That exact version of the event would not

⁶ A full-fledged counterfactual account will also claim that making a difference to the occurrence of an event is sufficient for causation. The point I want to make here merely concerns the claim that this is a necessary condition, which creates the problem from over-determination.

have occurred, and in that technical sense, R #1's decision made a difference. This sort of counterfactual dependence may be all that is required by the counterfactual dependence condition for causation.⁷

Allowing for versions of events and making them "fragile" or "fine-grained" may help us establish that both reviewers were causally involved in this article's rejection, even if we believe that some sort of counterfactual dependence is necessary for causation. This may be useful, e.g. if we want to handle over-determined harms with a causal non-comparative notion of harm.

However, even if some version of Lewis' strategy works, it will not save the counterfactual comparative notion of harm. I would not have been better off if any the versions of 'this article's rejection' where either reviewer #1 or reviewer #2 suggests acceptance had occurred, rather than the actual version. The over-determination challenge to the counterfactual dependence condition for causation should be kept apart from the over-determination challenge to the counterfactual comparative analysis of harm. Even if a plausible notion of causation admits that each reviewer decision is part of the causal history of the rejection, it is still the case that no single reviewer decision makes a difference to the rejection of the paper.

3 Diagnoses

Let me repeat this paper's simple main point. The plural harm principle says that *a plurality of events harms A if and only if that plurality is the smallest plurality of events such that, if none of them had occurred, A would have been better off*. But from the fact that a plurality of events is the smallest plurality of events such that *if none* of the events had occurred, A would have been better off, it does not follow that A would have been better off had the *plurality* not occurred. In the cases under consideration — the two rejection suggestions, X's and Y's shooting me, Batman's lethal heart attack and his simultaneously being hit by a flaming cannonball — the victim would *not* have been better off if the plurality of events picked out by the plural harm principle account had not occurred. This is because of the stipulation that in these cases, if any of the two elements in the plurality had not occurred, the other one would still have occurred, which implies that if the plurality consisting of the two events had not occurred, one of its elements would still have occurred. In these cases, this is sufficient for making the victim worse off to the same extent as when the plurality occurs. So, we must abandon either the plural harm principle or the view that making a difference to the worse for someone is essential to harming.

3.1 Closeness of Worlds and Contextual Variation

According to a standard way of thinking when we consider what would have happened, had not some actual event happened, the relevant counterfactual comparison is the possible world absent the actual event in question, which is closer to the actual world than any other possible world. A natural way of thinking might be to assume that the world without a certain *plurality* of events which is closest to the world where that plurality of events occurs, is one where another plurality of events occurs, which is minimally different from the actual plurality. If the

⁷ For criticisms and defenses of Lewis solution, see e.g. Schaffer 2001; Strevens 2003; Moore 2009, 410–425, Choi 2005; Petersson 2013.

actual plurality consists of more than one event, one might then think that the closest world is one with a plurality consisting in all but one of the members in the actual plurality. This would be consistent with what has been said so far about the plural harm principle as applied to the over-determination cases under consideration.

Admittedly, the question of how to precisify the vague notion of ‘closest world’ has no given answer, and on most views, this is a purely contextual matter. Michael McDermott makes the following observation.

Consider two unrelated events, e_1 and e_2 . Then (1) seems to be true, on its most natural interpretation, and (2) seems to be false.

- (1) If e_1 had not occurred, e_2 would still have occurred.
- (2) If they had not both occurred, one would still have occurred.

(Contrary to (2), we think that if they had not both occurred, that might have been because neither occurred.) (McDermott 2003, 227)

Intuitively McDermott’s observation seems plausible, at least if one tries hard to disregard (1) when considering (2) and vice versa.⁸ “What would have happened if one of the two events had not occurred?” and “What would have happened if both events had not occurred?” are questions we typically ask in different types of contexts. Typically, but not necessarily, the second question invites us to compare with a world where neither of the two events occur. In a similar manner, “What would have happened if this plurality of events had not occurred?” may invite us to compare with a world where none of the events in that plurality occur. I suggest that one possible explanation of why we find the plural harm principle intuitively attractive is that the formulation of the principle invites us to switch from thinking about the essential features of the original case, to thinking of a case where the relevant comparison is a world where none of the elements in the plurality occurs.

If we read the principle in that way, it is no longer applicable to the case that it is designed to handle. An essential feature of that case is that if any of the two events had not occurred, the other event would still have occurred, which implies that if the plurality had not occurred, one of its elements had still occurred. So, when we ask what would have happened if the plurality of R#1’s and R#2’s acts had not occurred, the context and hence the relevant counterfactual comparison is given by the case-description. There is no room for contextual variation in this respect between our reading of the description of the case and our reading of the principle when applied to that very case.

3.2 Distributive and Collective Readings of ‘Plurality’

In many real-life cases, pluralities of events fulfilling the plural harm principle do make a difference to the worse for someone. If we think of a murder committed by two accomplices, it is often reasonable to assume that if both had not acted the way they did, none of them had. We may think of the events in the plurality as co-dependent. Some of Feit’s and Parfit’s characterizations of the plural harm approach invites that kind of reading of the principle. They both think that it is plausible to say that in cases like the ones under consideration, the

⁸ McDermott suggests, but later dismisses, a way of making the conflicting verdicts on (1) and (2) compatible with Lewis’s analysis of counterfactuals, by assuming that the different judgments depend on different resolutions of the vagueness of ‘closer than’ in two different contexts. (McDermott 2003, 228)

two perpetrators ‘together harm me’ (Parfit 1984, 71, Feit 2015, 369). The term “together” carries with it a flavour of togetherness or collectivity, suggesting intentional co-ordination, planning, we-thinking or interdependence. When we are tempted to think in such terms of what would have happened if R#1 and R#2 had not rejected, or if A and B had not shot me, our intuitions may be misled. In the cases that the principle is supposed to resolve, the relevant events are *not* co-dependent.

The plural harm approach involves a shift from referring to the relevant events in plural, to referring to them with a noun in singular. ‘They harm me’, with a plural verb form and without further context, is ambiguous between a collective and a distributive meaning. It could mean that each does something to me, which is a natural reading when we think of two independent events like the two rejection suggestions, or the two shootings in Parfit’s example. It could also mean that the collection does something to me, which makes it natural to think of the events as co-dependent, or of each event as a necessary contribution to what happens to me.⁹ The singular verb form in ‘a plurality of events harms me’ points to a collective reading. Regardless of which noun we prefer – ‘group’, ‘set’, ‘whole’, ‘plurality’, ‘composite event’ – this kind of rephrasing tends to favour the collective reading.

As long as we do not try to disambiguate between the distributive and the collective reading, we may be able to hold on to the intuition that there must be some sense in which ‘they’ harm me in the rejection case. However, the interpretation that each harms me is incompatible with the counterfactual comparative analysis of harm. The interpretation that the collective harms me suggests that the events are co-dependent or that each event is necessary and together they are sufficient for my being harmed. Both features are incompatible with the over-determination case under consideration. So, a second tentative diagnosis of our being attracted by the plural harm account is that the formulations of the principle invites us to think of the cases in more strongly collective terms than the initial case-descriptions allow.

Feit’s explication of the term ‘plurality’ strengthens this suspicion. On the one hand, he says that ‘a plurality of events harms Batman’ is just as a way of saying that ‘they’ harm Batman, which is open between the distributive and the collective reading. On the other hand he says that the idea about plurality ‘is simply that some basic facts about harm are irreducibly plural, in the same way that the fact, say, that Alan and Beth and Carol and Don carried the boat might be irreducibly plural.’ (2015, 370) But unlike Batman’s heart attack and the flaming cannonball, four persons carrying a boat together appears to be a paradigm case of a collective activity, where typically each contribution makes a difference, co-ordination is required, etc.¹⁰

⁹ As Johansson & Risberg points out, when we say that events together harm a person, we typically “want to emphasize that none of the events would have harmed the person in the absence of the other ones — even if each of the events is such that, if it had not occurred, then the individual would have been better off”. But the plural harm principle is not designed to handle such cases. (Johansson and Risberg 2017)

¹⁰ A possibility that Feit dismisses would be to understand a plurality of events as a composite event. Feit thinks that this would commit him to a controversial metaphysical claim ‘that there exists some big spread-out event that consists of the heart attack and the flight of the cannonball’ in the Batman case. (2015, 370) I do not share that worry. We are free to individuate events in the manner that is most useful for our purposes, and that does not necessarily commit us to any controversial metaphysical views about what exists. A bookmaker would have no problem in accepting my bet on the event ‘the Vancouver Canucks reaches the finals and the UK wins the Eurovision Song Contest’, where the composite event is nothing more than the mereological sum of two unrelated events. This would not solve the problem I have pointed to, since the composite event ‘R#1’s suggesting rejection and R#2’s suggesting rejection’ makes no difference to my wellbeing in the over-determination case where one of the elementary events would have occurred if the composite event had not occurred. But I think that it might be an understanding of ‘plurality of events’ that is explicitly neutral with respect to degree of collectivity and interdependence.

3.3 The Smallest Actual Plurality and the Smallest Possible Plurality

A third feature of the analysis of plural harm that may mislead us to think of the pluralities in question as making a difference to the worse for someone is the necessary move to avoid including superfluous events in the plurality that harms – the restriction to the *smallest* plurality of events such that if none of them had occurred, A would not have been worse off. This is a restriction to the smallest *actual* plurality of events such that if none of them had occurred, A would have been better off. This may be conflated with the smallest *possible* plurality of events such that if none of them had occurred, A would have been better off, i.e. the plurality that would be sufficient and necessary for A's being worse off than A would have been if the plurality had not occurred. But in the cases under consideration the plurality is not necessary for A's being worse off in that sense since any of its elements would suffice.

4 Harmlessness, Wrongdoing and Responsibility

So, I suggest that we hold on to the counterfactual comparative analysis of harm, and consequently abandon the plural harm principle. Then we should conclude that neither Y's shooting me, X's shooting me, or the plurality of events consisting in their shooting me, harms anyone. Our conclusions should be similar with respect to the two rejections and Batman's two simultaneous causes of death. The consequences of taking this position are less radical than they might seem.

Accepting the implications of the counterfactual comparative analysis of harm is quite compatible with holding on to the part of Parfit's view which says that it can be morally *wrong* to perform an act when it belongs to the smallest plurality of acts such that if none of the acts had been performed, some person would not have been worse off. We could also claim that performing such an act can make the agent *co-responsible* for making the person worse off. Similarly, we could still say that an event can be *bad* if it belongs to the smallest plurality of events such that if none of the events had occurred, some person would not have been worse off.

Then, of course, we need an explanation of what makes the act wrong or the event bad, which does not appeal to the view that the act/event harms someone, or to the view that the act/event belongs to a plurality of events that harms someone. Moreover, we need an account of co-responsibility or complicity for an event which does not require that the co-responsible agent has made a difference to the occurrence of that event. These two issues should be separated from the conceptual analysis of 'harm' and they are strictly speaking beyond the scope of this paper. Here I will just give some hints about how I think they should be approached.

As Bradley points out, outside the act-consequentialist framework, there may be various explanations of how acts can be wrong even if they do not harm – explanations appealing to what a virtuous person would do, or to deontological side-restrictions, for instance. (Bradley 2012, 411) Even within a broadly consequentialist framework, it may be possible to find reasons for disvaluing events due to their belonging to some smallest set of events such that if none of the events had occurred, some person would not have been worse off. The strategy I suggest is to focus the facts about an event that qualifies it for membership in the relevant plurality in the first place. Those facts are facts about the effects of the elementary event – not about the effects of the plurality of events, nor

about the effects the event would have had if some other event had not occurred, but about the actual causal tendencies the elementary event initiates. Such causal facts about the event are what make it a member in some smallest set of events such that if none of the events had occurred, some person would not have been worse off.¹¹

When it comes to the question of co-responsibility or accountability, over-determination cases create apparent problems for what Christopher Kutz calls “The Individual Difference Principle” for accountability: “I am accountable for a harm only if what I have done made a difference to that harm’s occurrence.” (Kutz 2000, 125) The difference-making criterion for harm implies that X’s and Y’s acts are harmless, and the individual difference principle for accountability implies that neither X nor Y can be held to account for my death.

Kutz’s solution is to abandon the requirement of causal involvement for responsibility when we talk about collectively produced effects, and instead merely focus on “participatory intentions” in such cases. For reasons that I will not go into here, I think that this move reaches too far, and that it has unreasonable consequences. It is also unnecessary. We can reject the individual difference principle for accountability without giving up the requirement of causal involvement for accountability. Any reasonable theory of causation must be able to accommodate over-determination and pre-emption cases of the kind under consideration here. Some, like Michael S. Moore, think that this requires us to abandon the view that counterfactual dependence is a condition for causation. Others prefer to refine the counterfactual dependence condition for causation along the lines suggested by Lewis. Either way, it is clear that we can give up the individual difference principle for accountability without giving up the standard assumption that causal involvement is one of the pillars of moral responsibility.

None of this is incompatible with thinking that things harm a person only if they make things worse for that person than they would otherwise have been.

Acknowledgements Versions of this text were presented at the Higher Seminars in Practical Philosophy in Lund and in Uppsala, and at the Collectivity & Responsibility workshop in Lund 2018. I am grateful to participants in these events for useful and challenging comments. Special thanks to Wlodek Rabinowicz, Gunnar Björnsson, Caroline Touborg, Johannes Himmelreich, Jens Johansson, Olle Risberg, Per Algander, Olle Blomberg, Mattias Gunnemyr, and two anonymous reviewers. The work was supported by Grant P14-0212:1 from Riksbankens Jubileumsfond.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Bradley B (2012) Doing away with harm. *Philos Phenomenol Res* 85:390–412
 Choi S (2005) Understanding the influence theory of causation: a critique of Strevens. *Erkenntnis* 63(1):101–118
 Feit N (2015) Plural Harm. *Philos Phenomenol Res* 90:361–388
 Hanser M (2008) The metaphysics of harm. *Philos Phenomenol Res* 77:421–450
 Johansson J, Risberg O (2017) The preemption problem. *Philos Stud*. <https://doi.org/10.1007/s11098-017-1019-x>
 Kutz C (2000) *Complicity*. Cambridge U P, Cambridge

¹¹ The challenge for this strategy is to explicate the non-statistical teleological notion of a single case causal tendency towards realizing a certain event – “C tends to cause E” - in a metaphysically innocent way. In (2004, section IV), I suggest a way of understanding this concept.

- Lewis D (2000) Causation as influence, *Journal of Philosophy* XCVII
- McDermott M (2003) Closeness of worlds. *Acta Analyt* 18(30–31):227–230
- Moore MS (2009) Causation and responsibility, an essay in law, morals, and metaphysics. Oxford U. P, Oxford
- Norcross A (2005) Harming in context. *Philos Stud* 123:149–173
- Parfit D (1984) *Reasons and persons*. Oxford U P, Oxford
- Petersson B (2004) The second mistake in moral mathematics is not about the worth of mere participation. *Utilitas* 16(3):288–315
- Petersson B (2013) Co-responsibility and causal involvement. *Philosophia* 41(3):847–866
- Schaffer J (2001) Causation, influence, and effluence. *Analysis* 61(1):11–19
- Strevens M (2003) Against Lewis's new theory of causation: a story with three morals. *Pac Philos Q* 84:398–412