

Fragility, Influence, and Intrinsicity

PATRICK WILLIAMSON

AUSTRALIAN NATIONAL UNIVERSITY

Abstract

On David Lewis' original theory of causation, one event c causes another event e iff the right kind of counterfactuals are true. I outline the original counterfactual analysis before considering the challenge from redundant causation. I note that we might respond to redundant causation by adopting extreme standards of fragility, and that Lewis' criticisms of fragility are not persuasive. Lewis himself responds to redundant causation by developing an influence account of causation. I argue that influence is too imprecise as an analysis of causation, and that causation and causal influence in any case seem to be two distinct phenomena. I conclude by emphasising that the correct account of causation must allow for the intrinsicity of causal processes. This gives us good grounds for adopting Lewis' early account of causation as quasi-dependence, regardless of whether or not we view events as fragile. Lewis levelled trenchant criticisms against quasi-dependence, and I conclude by pinpointing the concessions that will, I expect, follow for advocates of quasi-dependence like myself.

1. Possible Worlds and Counterfactuals

A *possible world* is a causally-closed, maximal description of the way things could have been.¹ There is the set of all possible worlds, and this set includes our actual world, since our actual world is also a possible world.

Our world resembles some possible worlds more than others. We represent this fact with an *overall comparative similarity* relation across possible worlds. We measure similarity along two dimensions, both perfect match in particular matters of spatiotemporal fact and commonality in the laws of nature, but the relation will be inevitably vague inasmuch as it relies on our own subjective and context-dependent notions of resemblance.² By definition, every possible world falls somewhere in the similarity order and our own world most resembles actuality.³

This account of possible worlds proves pivotal in our analysis of counterfactual statements. Counterfactuals are subjunctive conditional statements, statements like *if I had been a banker I would have been rich*. The conversational implicature of these

¹ I assume this definition of possible worlds, putting aside complicating factors like modal realism.

² Lewis 1973, p. 560.

³ Ibid.

statements is a false antecedent – that I am not in fact a banker, even if I would like to be rich. We write these counterfactuals $A \square \rightarrow B$, *if A then would have B*.

Lewis gives the truth conditions for counterfactuals by employing possible worlds and the overall comparative similarity relation.⁴ A counterfactual is true if either the antecedent A fails in every possible world, or if some possible A and B world is closer to actuality than any A and $\sim B$ world.⁵ So the counterfactual *if I had been a banker I would have been rich* is true either if there are no possible worlds in which I am a banker, or if some possible world in which I am a rich banker is closer to actuality than any possible world in which I am a poor banker.

2. Counterfactual Causation

Lewis uses this account of possible worlds and counterfactuals to provide a counterfactual analysis of causation: *e* causally depends on *c* iff *had not c, not e*.⁶ More precisely, *e* causally depends on *c* if the right kind of counterfactual statements are true, both *had O(c), then O(e)* and *had $\sim O(c)$, then $\sim O(e)$* , where *O* indicates an event-tracking proposition.⁷ All this is given that *c* and *e* are distinct, non-overlapping, determinate events.⁸

We can illustrate this with an example. I am a secret agent, going about my business. I follow the target to an abandoned warehouse, pull my gun, light up the empty building with the flash of gunfire. The target sighs before collapsing on the concrete. Now, according to Lewis, the target's death causally depends on my shot if two counterfactuals are true, both *had I shot target would have died* and *had I not shot target would not have died*. Since the first proposition is true in the actual world, we consider the second proposition. If the second proposition is true – namely, if some *not shoot not die* world is closer than any *not shoot still die* world – then target's death causally depends on my shot.

We have in fact described *causal dependence*, but Lewis draws a distinction between causation and mere causal dependence.⁹ Causation supervenes on chains of causal dependence: if *c* causally depends on *b*, and *b* causally depends on *a*, then *a* causes *c*. This allows us to maintain that causation is transitive, even though causal dependence may or may not be transitive.¹⁰

So, to revisit our example: I cause the target's death so long as his death causally depends on the bullet, which causally depends on my pistol barrel, which causally depends on my trigger finger. I cause the death, in other words, when there is a chain of causal dependence connecting my trigger finger to the death. This remains the case

⁴ Ibid.

⁵ Ibid.

⁶ Ibid., p. 562.

⁷ Ibid., p. 563.

⁸ For the purposes of this paper these are only stipulations. For a discussion of indeterminate counterfactual causation, see Lewis 1986.

⁹ Lewis 1973, p. 562.

¹⁰ Ibid., p. 563.

even if there is not *direct* causal dependence from my trigger finger to the target's death – after all, perhaps a second agent waits in the dark, ready to bring about the target's death in every possible world in which I fail to pull the trigger, hence destroying any direct counterfactual dependence from the target's death to my trigger finger.

3. Redundant Causation

I now turn to one of the strongest objections facing the original counterfactual analysis, the issue of redundancy. These are causal scenarios in which either of two events c_1 and c_2 would be sufficient to entail event e , but as a result, e counterfactually depends on neither.¹¹

I restrict myself to the most difficult type of redundancy, so-called late pre-emption cases. Here is an example from.¹² Billy and Suzy are out playing in the sun. They each chuck a rock at an old glass bottle. Suzy's rock soars through the air, crashes into the bottle, and explodes glass fragments across the dirt. Billy's rock passes through empty air where the bottle once stood and sails off into the distance.

Suzy is clearly the cause of the break, since it was her rock that crashed into the bottle. But the counterfactual analysis tells us that Suzy is *not* a cause of the break. After all, if Suzy's rock hadn't broken the bottle, then Billy's rock would have. So the counterfactual *had Suzy not thrown, the bottle would not have broken* is false. So the counterfactual analysis says that Suzy's throw is non-causal.

As a rule, we see that the plurality of potential causes in redundancy cases prevents the right counterfactuals from emerging. It is simply not true that *had not c_1 , not e* , since then c_2 would have entailed e . It is simply not true that *had not c_2 , not e* , since then c_1 would have entailed e . But without these counterfactuals, the counterfactual analysis cannot diagnose the cause of e . So we must either revise or abandon the counterfactual theory of causation.

4. The Fragility Response

We might respond to these late pre-emption cases by adopting extreme standards of fragility. When we say that an event e is *fragile*, we say that e would have been some different event altogether under some miniscule variation in spatiotemporal fact.

By adopting extreme standards of fragility, we rule out troublesome cases of redundant causation *a priori*. Why? It is simply not possible – in our world, at least – for any two events c_1 and c_2 to independently bring about an identical event e . Rather, they bring about two distinct fragile events e_1 and e_2 .

So the analysis goes as follows. Suzy and Billy each chuck a rock at the bottle. Each throw would have led to a different fragile event upon collision, *Suzy's Break* and

¹¹ Lewis 1986, p. 193.

¹² Lewis 2000, p. 82.

Billy's Break. Hence we have a perfectly respectable chain of counterfactual dependence, *Suzy's Throw* $\square \rightarrow$ *Suzy's Break*. Hence the standard counterfactual analysis gives the correct answer, that *Suzy's Break* counterfactually depends on *Suzy's Throw*. Hence *Suzy's Throw* causes *Suzy's Break*.

Lewis offers two arguments against adopting fragility. But in my mind fragility remains a viable response to the problem of pre-emption. In the following section I try to show why Lewis' criticisms are unpersuasive.

First, Lewis does not think extreme fragility fits with our folk intuitions regarding eventhood.¹³ We usually allow that an event might have varied in spatiotemporal minutiae without being instantly destroyed and replaced by another event altogether. For instance, the concert could have been postponed but remained the same concert, or the seminar could have been postponed while remaining the same seminar.¹⁴

This reason strikes me as insufficient. First, we can grant Lewis that our metaphysical analysis should (in broad and general terms) track our folk intuitions. But this fact alone cannot be used as a critique of any single philosophical doctrine without elevating folk intuitions to a status with which many of us will be uncomfortable: Lewis' reply gives the impression that folk intuitions are the primary grounds on which a metaphysical theory is to be accepted or rejected.¹⁵ Perhaps things would have been different if Lewis merely held up folk intuitions as an adjudicatory tool when assessing equally meritorious metaphysical theories.¹⁶ Second, all of this has assumed that there *is* some folk intuition when it comes to the fragility or non-fragility of events. But this point also seems open to dispute. It strikes me that 'the folk' refer to the delayed concert as the 'same' concert as a matter of convenience, without giving much thought to the underlying metaphysical status of events themselves. The conductor is interested in whether the musicians are ready, whether the house is full, whether his suit is straight and tidy – in all of these things – but *not* in whether the concert would have maintained its essential metaphysical identity if any of these individual components were to vary across space and time. After all, how do we know what the conductor says when cornered by a pack of philosophers? Perhaps the conductor gives us a nervous glance, admits that, *technically* speaking, the delayed concert would have been a different event, before adding that it is not very useful to get caught up in such philosophical technicalities. In sum: we should not assume that there is some established folk consensus on the fragility or non-fragility of eventhood.

Second, Lewis argues that extreme fragility allows too many events to qualify as spurious causes.¹⁷ There are hundreds of events that go into making *Suzy's Break* the perfectly fragile event that it is, including air pressure, rock-type, glass-thickness, wind-currents, and so on. If any of these factors had failed to obtain, some other fragile event would have obtained in the place of *Suzy's Break*. Hence, *Suzy's Break*

¹³ *Ibid*, p. 86.

¹⁴ *Ibid*.

¹⁵ *Ibid*.

¹⁶ Thank you to an anonymous referee.

¹⁷ Lewis 1986, p. 198.

counterfactually depends on all of the aforementioned background details. Hence, all these background details are spurious causes for *Suzy's Break*.

Lewis is right to raise the problem of spurious causes. But he is wrong to level it as a probative objection against fragility. Why? The standard counterfactual analysis was already infected by the problem of spurious causation. On the standard counterfactual analysis my smoking causes my lung cancer, since *had I not smoked I would not have had lung cancer*. But on the standard counterfactual analysis my lungs also cause lung cancer, since *had I not had lungs I would not have had lung cancer*.¹⁸ So in buying into the standard counterfactual analysis, we have already allowed for spurious causation.

Now, Lewis might object that the range of spurious causes permitted by the original counterfactual analysis is relatively slim, but the range of spurious causes permitted by the fragility account is virtually infinite. But this response seems incorrect so long as we maintain that causation is transitive. If causation is transitive then my parents cause my birth, and my birth causes my lungs, and my lungs cause my cancer; ergo, my parents cause my lung cancer. And the same might be said for my grandparents, great-grandparents, and so on, back down the causal chains of history. So it is not obvious that the original counterfactual analysis allows only a restricted range of spurious causes. What Lewis would have to argue is that some types of spurious causes are more repulsive than others – that fragility allows a certain *type* of spurious cause not previously allowed by the counterfactual analysis. And I maintain that all spurious causes are equally spurious – it does not matter whether my grandparents are a spurious (counterfactual) cause for my lung cancer, or whether the background breeze is a spurious (fragile) cause for *Suzy's Break*.

5. Influence

Lewis later responded to late pre-emption cases by recasting causation in terms of *influence*. On the influence account, one event *c* causes another event *e* if there is a "chain of stepwise influence from *c* to *e*".¹⁹ There is causal influence from *c* to *e* when a range of close alterations to *c* correspond to sufficient distinct alterations of *e*. So, influence emerges when there is counterfactual dependence of "whether, when and how on whether, when and how".²⁰ Since we are considering a wide range of possible alterations, influence is measured in degrees rather than absolutes.²¹

The influence account makes neat work of late pre-emption cases. *Suzy's* throw causes the break since it exercises a high degree of influence over the break. A little faster and the bottle falls on its side; a little harder and the hole takes a different shape; and so

¹⁸ Menzies 2004, p. 143.

¹⁹ Lewis 2000, p. 91.

²⁰ Ibid.

²¹ Some examples may prove useful. On the influence account, the following scenarios all qualify as causal: my waking up in the morning *whether* depends on *whether* my heart continues to function, *when* depends on *when* my alarm clock goes off, *how* depends on *how* many hours I slept, *how* depends on *whether* my friend throws a bucket of water on me, and so on. The original counterfactual analysis only adjudicated causation on the basis of *whether / whether* dependence, for instance, *whether* I wake up depends on *whether* my heart continues to function.

on.²² By contrast, Billy's throw fails to cause the break since he exercises a low degree of influence over the break. This is not to say Billy exercises no influence, but only that he exercises comparatively less than Suzy.

One other example, Lewis' trumping case, will prove useful in the following section.²³ Sergeant and Major both shout advance and the troops march forward. The troops follow Major's order, since he is the commanding officer. But either shout would have sufficed, giving that Sergeant and Major are both commanding officers. Now, who causes the advance? The influence account tells us, correctly, it seems, that Major is the cause. After all, the right range of alterations on his command c produce significant alterations in the troops' advance e . Since Sergeant is outranked, variations on his shout produce fairly little variation in the soldiers' conduct.

6. Against Influence: Imprecision

It seems to me that there are two good reasons for rejecting the influence account of causation. Here is one reason: influence admits of too much impreciseness and context-dependency to be a viable theory of causation.

First. How large is the set of possible alterations to c ? Unfortunately there is no precise answer. We want 'close' alterations to c , but 'close' is a subjective notion that varies from conversational context to conversational context. There is no clear answer, for instance, as to whether we consider worlds in which (a) the Major shouts the same word slightly differently, (b) the Major shouts different words, (c) does not shout at all, or (d) performs some other action instead of shouting.

Second. Suppose alteration c^* results in alteration e^* . When is e^* sufficiently 'different' from e for this pair $c^* \square \rightarrow e^*$ to contribute to the influence of c on e ? Is it enough if e and e^* are one atom apart? Or must e and e^* be thoroughly and unmistakably different? Again, there is no precise answer. We cannot simply count atoms, because we are also measuring *qualitative* differences, for example, differences in the manner of e . But the problem seems worse this time round. I don't know how to judge whether or not an alteration e^* is 'different' enough without relying on our preconceived notions of causation – in other words, without asking if the change from e to e^* is significant enough to qualify the previous event c as a *cause*. But what qualifies as causation is precisely the question at hand. So I don't know how to determine whether an alteration in e is sufficiently 'different' in any way that is non-circular.

Third. Whether or not c has enough causal influence to qualify as a cause of e depends on the other causal candidates for e under consideration. But the number of other causal candidates under consideration is, again, an imprecise matter determined by conversational context. I take this to be evident in the case of Sergeant versus Major. Sergeant *does* exercise a substantial amount of causal influence over the soldiers (had he shouted a second earlier the troops would have marched, had he shouted a little

²² Ibid., p. 92.

²³ Ibid., p. 81.

louder he would have blotted out Major's voice, and so on). But we deny that Sergeant is the cause because influence claims are comparative, and given that we are comparing him to Major in this particular conversational context, he does not exercise enough causal influence to qualify as a cause.

So the influence account is infected with at least three degrees of imprecision – how large the range of alterations to c , how to know whether a variation from e to e^* is significant enough to contribute to influence, and the range of other causal influencers under consideration. As I mentioned, some of these imprecisions hide a certain circularity. But in general, my concern is that a theory of causation with this degree of imprecision cannot tell us, in any objective sense, what is causation and what is not. As to whether any x causes e , there is no clear answer. We can only give an account of how we might talk about x causing e , conditional on vast quantities of imprecise information being filled in by the background conversational context. The original counterfactual analysis was not free from such imprecision, but limited imprecision to the overall comparative similarity relation between possible worlds.²⁴

7. Against Influence: Coming Apart

Here is another reason for rejecting influence. It seems to me that causation and causal influence are two distinct phenomena that might come apart. Genuinely causal relations might exhibit varying degrees of causal influence – a kind of causal hallmark – but we should not make the mistake of conflating the two.

Here is an initial thought to lead us in this direction. Suppose we reconsider the case of Sergeant and Major, but radically simplify the scenario such that Major is only capable of doing two things – barking ‘advance’ and barking ‘stop’. Perhaps this comes from a lifetime of mind-numbing drill training, and if Major attempts to say anything else he will simply collapse on the asphalt in a fit of cognitive dissonance, the troops letting out a familiar sigh. It seems that we should now say two things. First, that Major's command is still the cause of the advance – after all, the troops are listening to Major and not Sergeant. Second, that Major's command exercises *very little* causal influence over the troops. Alterations in his shout only result in a collapsed officer flat on the parade ground. So, by slightly modifying the scenario we find that causation is present but influence is absent.

This coming apart of causation and causal influence becomes clear in other less-contrived counterexamples. Take the following example adapted from Strevens.²⁵ Billy and Suzy have gone to the pub, and are now throwing darts at a dartboard. Billy has cheated and weighted Suzy's dart so that it flies low. Just as Suzy throws her dart, a wind blows through the window, and causes it to fly just as far right as it flies low:

²⁴ Lewis 1973, p. 560.

²⁵ Strevens 2003, pp. 403–04.




fig 1. weighted dart causes shot to fly low. wind causes shot to fly right.

Now, what causes Suzy's dart to fly low? A range of variations on Billy's weight b produce a range of fine-grained variations on the shot's landing low. But a range of variations on the wind w produce just as many fine-grained variations on the shot landing low. So, by influence, the horizontal wind is just as much a cause for flying low as the weighted dart is a cause for flying low.

Someone might object: *the wind doesn't influence the lowness, since the dart will always fly low no matter how hard the wind blows!* But this is to reject Lewis' conception of influence. On Lewis' view the wind must influence the lowness of the dart exactly as much as Billy's weight does, because the same degree of spatial alterations on 'lowness' can be produced by altering either cause.²⁶

In the dart case we are inclined to say that the wind influences the lowness but is not the cause of the lowness. In the modified Major case we are inclined to say the reverse, that Major is a cause of the advance but not an influencer of the advance. But if causation and influence can come apart in both directions, this is good evidence that causation and influence are two distinct phenomena. So distinct, in fact, that we should not ground our account of one in the other.

8. Intrinsic Processes

I have criticised the influence account of causation. But in light of these counterexamples we can say more about the nature of causation itself, a point that may not have been explicit in the mere fragility account. Suzy's bottle, the shouting Major, and the weighted dart all have something in common: in each case, the 'correct' cause is connected to the effect via some intrinsic series of spatiotemporal events. For instance, we can trace the rock as it leaves Suzy's fingers, slips through the air, and crashes into the glass. We can trace the Major's shout as it echoes across the training field, into the soldiers' eardrums, and causes certain neurons to fire. We can trace the weight as it angles the dart down towards the floor, right until the metal tip sinks into the corkboard.

In light of these examples we can posit the following: the correct analysis of causation will respect the intrinsicity of causation, that is, it will adjudicate causation based on the direct series of spatiotemporal events connecting cause and effect in and of themselves.

²⁶ Also see *ibid.*, p. 404.

To this end we can employ Lewis' early account of causation as *quasi-dependence*.²⁷ On the quasi-dependence account, *c* causes *e* if there is either a chain of causal dependence or quasi-dependence from *c* to *e*. *e* quasi-depends on *c* if *e* counterfactually depends on *c*, not in the actual world, but in some possible world in which both (1) the same intrinsic series of events connects *c* to *e*, but (2) the extrinsic background details of the causal scenario vary, extrinsic background details including pre-emptors.

So we take Suzy's throw; we go to some possible world where the same intrinsic process connects the throw to the break; we remove extraneous background details like Billy's throw; we find that the right type of counterfactual dependence arises; hence the break quasi-depends on her throw, and her throw is the cause of the break.

One key advantage of the quasi-dependence account is that it captures the intrinsicity of causation, while still allowing a certain agnosticism about whether events are fragile or non-fragile. We do not need to know whether the bottle break is fragile in order to determine its cause. We only need to know whether the right counterfactuals arise in intrinsically identical cases. For this reason, I now defend a quasi-dependence analysis of causation, putting aside the issue of fragility as a secondary theoretical consideration.

9. Concessions for Quasi-Dependence

Lewis later offered a range of reasons for rejecting quasi-dependence (QD) and the intrinsicity of causation. These are not trivial criticisms, and deserve an extended treatment of their own. But I want to close by getting clear about what type of concessions I think the defender of QD must be willing to make in light of Lewis' critique. I restrict myself to Lewis' two strongest criticisms.

QD cannot handle trumping cases, so says Lewis.²⁸ In trumping cases we have one cause, Major. And yet we have *two* completed intrinsic event chains leading to the effect, one from Major and one from Sergeant. An implicit consequence: if we attempt to analyse trumping cases via QD, we will be forced to admit that both trumper and trumpee are causes.

I suspect the defender of QD must maintain that all trumping cases are cutting cases.²⁹ That is, the defender of QD must maintain that all trumping cases involve one and only one completed causal chain – a chain which blocks any other potential causal chains from running to completion. The Major's command echoes across the yard and into the soldiers' eardrums, causes certain neural pathways to fire, and at some point these new pathways cut certain Sergeant-obedient neural pathways that would have otherwise run to completion. Hence, Major severs the intrinsic causal chain connecting Sergeant to the advance. So there is only one intrinsic cause-effect chain, and QD can still give a satisfactory analysis of the case.

²⁷ Lewis 1986, p. 206.

²⁸ Lewis 2000, p. 83.

²⁹ *Ibid.*, p. 81.

Now, Lewis imagines alien possible worlds in which there *are* non-cutting trumping cases.³⁰ Perhaps Merlin trumps Morgana when it comes to casting a spell of transmutation. My concern with this answer is methodological: when constructing a theory of causation, we should first see whether our theory maps the actual world, nomologically accessible worlds, and nomologically approximate worlds, since these are the worlds in which our causal intuitions and our causal reasoning remain the sharpest. If our causal theory successfully maps these worlds, then we should see whether our theory maps alien worlds. The fact that a theory does not map alien worlds cannot be used as an objection to a causal theory until its competitors are shown to perform both better in the actual world, nomologically accessible, and nomologically approximate worlds, as well as in alien worlds containing, say, Merlin and Morgana.

QD cannot handle double prevention, so says Lewis.³¹ These are cases in which I cause an outcome by preventing a preventor – by preventing the President from firing nuclear missiles, I prevent Russia from retaliating, and by preventing Russia from retaliating I cause Joe to have his breakfast like normal.³²

Why do double prevention cases pose a challenge to QD? Because these cases are causal, but crucially, because they do not instantiate a chain of spatiotemporal events connecting cause to effect. Such a chain would be required for the intrinsicity of causation. So QD fails.³³

My first inclination was to try and find some intrinsic chain of events at work in double prevention cases. But how could there be, if omissions like *not pressing the button* are not concrete events in spatiotemporal reality? So I think the simplest solution is for defenders of QD to deny that omissions exist, and hence that omissions can be causes.

So Lewis presents a challenge to QD, an omissive causal string like *My Prevention* → *Omitted Button Press* → *Omitted Russian Missiles* → *Joe's Breakfast*. But the defender of QD denies that omissions exist. It follows that the string is some kind of misnomer or misdescription of reality. Hence it does not pose a challenge to the QD account of causation.

It seems to me that the defender of QD must deny the existence of absences altogether. They cannot maintain that absences merely exist in a non-causal way, because if absences exist then the absences will generate the right counterfactuals via their negations and hence qualify as causes on a counterfactual analysis.

³⁰ Ibid.

³¹ Ibid., pp. 83–84.

³² McDermott 1995, p. 529.

³³ Lewis 2000, p. 85. These cases do not instantiate a complete chain of spatiotemporal events connecting cause to effect, since omissions are serving as causes, and since omissions are non-spatiotemporally located.

I am not sure if it is radical to wholeheartedly deny the existence of omissions. The defender of QD will have to find other actual pieces of spatiotemporal reality to cause Joe's breakfast via the right chains of counterfactual dependence.

Instead of admitting absences into our ontology, perhaps we say something like the following: *the President omitted to press the button* is not the claim that a certain omission exists, but rather, a contrastive positive claim about reality, a shorthand for *the President was sitting down and combing their hair in the other room, far away from the button*. In this way, the defender of QD might redescribe omissive causal strings as non-ommissive causal strings, and hence, provide a causal analysis still grounded in unproblematic chains of intrinsic spatiotemporal process.^{34, 35}

So the defender of QD has been forced to make two concessions, that trumping cases are cutting cases and that omissions cannot be causes because they do not exist. I find both concessions plausible, and hence I find quasi-dependence a plausible theory of causation. I prefer it even to the fragility account, inasmuch as quasi-dependence foregrounds the essential intrinsicity of causation.

10. Conclusion

I began by considering Lewis' original counterfactual analysis of causation, arguing that we might reasonably respond to redundancy by adopting extreme standards of fragility. Lewis would rather resolve redundant late pre-emption by recasting his original counterfactual analysis in terms of influence, but I think this is mistaken, since influence is too imprecise as an analysis of causation and since there are cases in which influence and causation come apart. Regardless of whether or not events are fragile, the correct account of causation will respect intrinsicity, the fact that causation is somehow to do with chains of spatiotemporal events connecting cause to effect in and of themselves. To this end, Lewis' early account of causation as quasi-dependence is a profitable analysis of causation, although I suspect it requires certain non-trivial metaphysical concessions.³⁶

³⁴ Hart & Honore 1985, p. 38.

³⁵ If we deny omissions, do we undermine the entire counterfactual analysis so far, which depended on strange things like $\sim c$ and $\sim e$? Only if we think that $\sim c$ and $\sim e$ are claims about the existence of omissions: instead, as I have suggested, we might view $\sim c$ and $\sim e$ as contrastive positive claims about reality. c reads: c exists. $\sim c$ reads: x exists.

³⁶ Thank you to three anonymous referees for their helpful comments and suggestions.

References

- Hart, H. and Honore, A. (1985) *Causation in the law*. Oxford: Clarendon Press.
- Lewis, D. (1973) 'Causation', *The Journal of Philosophy*, 70:556–567.
- Lewis, D. (1986) 'Postscripts to causation', in *Philosophical Papers Vol. 2*, 172–213. Oxford University Press, Oxford.
- Lewis, D. (2000) 'Causation as influence', *The Journal of Philosophy*, 97:182–197.
- McDermott, M. (1995) 'Redundant causation', *British Journal for the Philosophy of Science*, 46:523–544.
- Menzies, P. (2004) 'Difference making in context', in J Collins, N Hall & LA Paul (eds) *Causation and Counterfactuals*, 139–180. MIT Press, Cambridge.
- Strevens, M (2003) 'Against Lewis's new theory of causation: a story with three morals', *Pacific Philosophical Quarterly*, 84:398–412.