Synthese
https://doi.org/10.1007/s11229-022-03596-z

ORIGINAL RESEARCH

Relativizing proportionality to a domain of events

Caroline Torpe Touborg[1]

Received: 19 July 2021 / Accepted: 25 January 2022
© The Author(s) 2022

Abstract
A cause is proportional to its effect when, roughly speaking, it is at the right level of detail. There is a lively debate about whether proportionality is a necessary condition for causation. One of the main arguments against a proportionality constraint on causation is that many ordinary and seemingly perfectly acceptable causal claims cite causes that are not proportional to their effects. In this paper, I suggest that proponents of a proportionality constraint can respond to this objection by developing an idea that is present in Yablo's early work on proportionality, but which has strangely been ignored by both Yablo and others in the subsequent debate. My suggestion is that proportionality—and, indeed, causation itself—is relative to a domain of events. At the metaphysical level, this means that the causal relation has an extra relatum—namely, a domain of events. At the level of language, it introduces a new way in which causal claims are context-sensitive: what is expressed by a causal claim depends on the contextually relevant domain of events. As I argue, this suggestion allows us to accommodate the truth of ordinary causal claims while extending the explanatory benefits of a proportionality constraint.

Keywords Causation · Proportionality · Yablo · Domain of events

1 Introduction

It is an open question whether proportionality is a necessary condition for causation. In this paper, I suggest that *if* you want to hold that proportionality is necessary for causation, you should understand proportionality in a particular way—namely, as being relative to a domain of events.

Roughly, a cause is *proportional* to its effect when it is at just the right level of detail relative to its effect. Suppose, for example, that the pigeon Sophie has been trained to peck at red to the exclusion of all other colours. She is presented with a scarlet object

✉ Caroline Torpe Touborg
caroline.touborg@gmail.com; caroline.touborg@umu.se

[1] Department of Historical, Philosophical and Religious Studies, Umeå University, Umeå, Sweden

Published online: 11 March 2022

and pecks at it. In this case, *the object's being red* is at just the right level of detail relative to Sophie's pecking, since Sophie is precisely responding to the redness of the object. By contrast, *the object's being scarlet* includes too much detail—it doesn't matter to Sophie whether the object is scarlet or some other shade of red; and *the object's being coloured* leaves out too much detail—it *does* matter to Sophie whether the object is red or some other colour. In this case, then, *the object's being red* is a proportional cause of Sophie's peck, while *the object's being scarlet* and *the object's being coloured* are not proportional.[1]

It is an open question whether proportionality is necessary for causation.[2] One of the main arguments against taking proportionality to be a necessary condition for causation is this: if proportionality is necessary for causation, then many ordinary causal claims that seem perfectly acceptable turn out to be false. For example, (1) below seems perfectly acceptable:

(1) Socrates' drinking hemlock caused his death.

However, (1) fails to satisfy the condition of proportionality, since it includes too much detail (Bontly, 2005, p. 340): Socrates would still have died if he had drunk some other lethal poison; by specifying the particular kind of poison, *Socrates' drinking hemlock* therefore includes too much detail, just as *the object's being scarlet* includes too much detail by specifying the precise shade of red.

The difficulty exemplified by (1) has been pointed out repeatedly in the literature on proportionality (see e.g. Bontly, 2005; Franklin-Hall, 2016; Maslen, 2017; McDonnell, 2017; Shapiro & Sober, 2012; cf. McGrath, 1998): while a proportionality constraint on causation may seem initially plausible, it would force us to reject a multitude of ordinary and seemingly perfectly acceptable causal claims as false and replace them with anaemic statements such as (2) or even (3):

(2) Socrates' drinking a lethal poison caused his death.
(3) Socrates' doing something fatal caused his death.

I agree that this is indeed an unacceptable consequence of taking proportionality as a necessary condition on causation. However, instead of simply conceding defeat, I suggest that proponents of a proportionality constraint on causation should go contextual: proportionality should be understood relative to a contextually determined domain of events. Indeed, Yablo suggests this kind of approach in his early work on proportionality (1992a), but neither he nor others have followed it up in the subsequent literature. I believe it has been a mistake to abandon this idea: as I show in the following, a contextual account of proportionality allows us to reconcile a proportionality constraint with the truth of (1), while extending the proportionality constraint's explanatory power.

I proceed as follows: first, I present the notion of proportionality in more detail (Sect. 2) and show how a simple proportionality constraint runs into trouble (Sect. 3). Next, I suggest, drawing on Yablo, that proportionality should be understood relative to a domain of events (Sect. 4). I show how such a relativized notion of proportionality

[1] The example is due to Yablo (1992a, p. 432; 1992b, p. 257).

[2] For example, List and Menzies (2009) hold that proportionality is necessary for causation; Bontly (2005), McDonnell (2017), Weslake (2013, 2017), and Woodward (2015, 2021) suggest that proportionality is merely a feature that some causal relations have and others lack.

can accommodate our intuitive verdicts in the trouble cases (Sect. 5) and that it can do important explanatory work (Sect. 6). I end with a discussion of the considerations that may go into choosing a domain of events (Sect. 7).

2 Proportionality

The notion of proportionality was originally developed by Yablo (1992a, 1992b). Yablo's account is often taken as the starting point for discussions of proportionality (see e.g. Bernstein, 2014; Maslen, 2017; McDonnell, 2018; Sartorio, 2010), and I shall also do so in this paper. In this section, I present a pared-down version of Yablo's account. While I shall relate my discussion to Yablo's account in the following, however, it is worth noting that both the problem and my suggested solution may apply to other accounts of proportionality as well.[3]

Yablo's account of proportionality has two parts. The first part concerns the causal relata. The second part concerns the conditions for proportionality itself.

Let us begin with the causal relata. To get the proportionality conditions up and running, we need to think of the causal relata as events, understood in a particular way: for any given space–time region, there is a multitude of events that all occur precisely within this space–time region, and which differ from each other only in terms of their modal profiles.[4] Consider, for example, the space–time region where Sophie is presented with the scarlet object.[5] There is a multitude of events that all occur precisely within this space–time region. One of these events is essentially the presentation of a scarlet object. We may call this event *scarlet*. Another of these events is essentially the presentation of a red object (it is only accidentally the presentation of a scarlet object). We may call this event *red*. A third event is essentially the presentation of a coloured object (it is only accidentally the presentation of a red one). We may call this event *coloured*.

The three events *scarlet*, *red*, and *coloured* are numerically distinct, because they have different conditions of occurrence. For example, *red* occurs in worlds where Sophie is presented with an object in some other shade of red—say crimson; *scarlet* does not. However, while the three events are numerically distinct, they stand in interesting logical relations to each other: *coloured's* conditions of occurrence are strictly less demanding than *red's*; *red's* conditions of occurrence are strictly less demanding than *scarlet's*. When one event's conditions of occurrence are strictly less demanding than another's, we shall say that the first event is a *weakening* of the second (and the second is a *strengthening* of the first):[6]

[3] More precisely, my proposed solution (cf. Sect. 4) works for any account of proportionality that involves quantification over events, such as e.g. Weslake (2013).

[4] For a well-known implementation of this picture, see e.g. Lewis (1986). Yablo's picture is more complex, but the differences do not matter here.

[5] Saying "*the* spacetime region" is of course not accurate, since there are many different and seemingly equally good ways to draw the relevant boundaries. However, nothing hinges on this; in the following, I shall therefore continue to speak of *the* spacetime region.

[6] Once again, Yablo's picture is more complex.

Weakening: x^- is a weakening of x (and x is a strengthening of x^-) just in case the conditions of occurrence of x^- are strictly less demanding than the conditions of occurrence of x.

Following Yablo, I shall use the notation "$x^- < x$" to mean that x^- is a weakening of x, and the notation "$x^+ > x$" to mean that x^+ is a strengthening of x.

This brings us to the conditions for proportionality. On Yablo's picture, events that are logically related in this way—where one event is a weakening (or strengthening) of the other—compete for the status as cause.[7] For example, the three events *scarlet*, *red*, and *coloured* compete with each other for the status of being a cause of Sophie's peck. The most proportional of the events in the competition wins and counts as a cause; the others do not count as causes.

What does it take to win the competition? To win, Yablo suggests, an event must satisfy two necessary and jointly sufficient conditions.[8] Each of these two conditions captures one side of what it is for a cause to include just the right amount of detail relative to its effect: the first condition captures the idea that a cause should not include *too much* detail relative to its effect; the second condition captures the idea that it should not include *too little*.

The first condition is that the winner should be *required* for the effect, where being required is defined as follows:

Required: x is *required* for y just in case for all $x^- < x$,
if x^- had occurred without x, then y would not have occurred.

This condition captures the idea that, compared with the other competitors, a cause should not include *too much* detail relative to its effect. In the case of Sophie, for example, *scarlet* fails to satisfy this condition because it is knocked out by *red*: *red* is a weakening of *scarlet*, and if *red* had occurred without *scarlet*—that is, if Sophie had been presented with an object that was red without being scarlet, e.g. a crimson object—then Sophie would still have pecked. Thus, considering *red* reveals that *scarlet* is not required for Sophie's peck. That is exactly the result we wanted since, intuitively, *scarlet* includes too much detail—its conditions of occurrence are too demanding—for it to be proportional to Sophie's peck.

The second condition is that the winner should be *enough* for the effect, where being enough is defined as follows:

Enough: x is *enough* for y just in case for all $x^+ > x$,
if x had occurred without x^+, then y would still have occurred.[9]

[7] In fact, Yablo limits the competition to determinables and their determinates. Following Bontly (2005, p. 335), I believe it is unnecessary to tie the notion of proportionality this closely to the determinable-determinate relation: any two events where one is a weakening of the other, as defined above, may compete for the status as cause.

[8] Yablo's conditions do not deliver intuitively correct verdicts about proportionality in cases of redundant causation, as Yablo himself is fully aware (see e.g. Yablo, 1992b, footnote 60; Yablo, 2003, footnote 16). In the following, we shall not be considering cases of redundant causation, and I therefore set this issue aside.

[9] According to Yablo's original condition, x is enough for y just in case for all $x^+ > x$, x^+ is not required for y. The simplified version that I use here is found in McGrath (1998), and later in e.g. Bontly (2005) and Maslen (2017). McGrath's condition is more demanding than Yablo's: whenever it is the case that if x had occurred without x^+, then y would still have occurred, it is also the case that x^+ is not required for y, but

This condition captures the idea that, compared with the other competitors, a cause should not include *too little* detail relative to its effect. In the case of Sophie, for example, *coloured* fails to satisfy this condition because it is knocked out by *red*: *red* is a strengthening of *coloured*, and if *coloured* had occurred without *red*—that is, if Sophie had been presented with an object that was coloured without being red, e.g. a blue object—then Sophie would not have pecked. Thus, considering *red* reveals that *coloured* is not enough for Sophie's peck. Again, that is exactly the result we wanted since, intuitively, *coloured* includes too little detail—its conditions of occurrence are not demanding enough—for it to be proportional to Sophie's peck.

In the end, *red* emerges as the winner, since it is both *required* and *enough* for Sophie's peck, while *scarlet* fails to be required and *coloured* fails to be enough. Yablo's verdict, therefore, is—as it should be—that *red* is proportional to Sophie's peck, while *scarlet* and *coloured* are not. To sum up:

Proportional: x is *proportional* to y just in case
(a) x is required for y, and
(b) x is enough for y.

As we shall see in the next section, however, trouble is not far to seek.

3 Trouble for proportionality

In the case of Sophie, the idea that proportionality is a necessary condition for causation accommodates our intuitive verdicts about the appropriateness of causal claims: choosing between the three claims below, it seems perfectly appropriate to assert (4), but inappropriate to assert (5) or (6):

(4) The object's being red caused Sophie to peck.
(5) The object's being scarlet caused Sophie to peck.
(6) The object's being coloured caused Sophie to peck.

On the natural assumption that "the object's being red" refers to the event *red*, "the object's being scarlet" refers to *scarlet*, and "the object's being coloured" refers to *coloured*, this perfectly matches the verdicts of Yablo's account.

However, as has been widely noted in the literature (see e.g. Bontly, 2005; Franklin-Hall, 2016; Maslen, 2017; McDonnell, 2017; Shapiro & Sober, 2012; cf. McGrath, 1998), many ordinary causal claims that seem perfectly appropriate fail to satisfy the condition of proportionality. We have already seen an example in the introduction:

(1) Socrates' drinking hemlock caused his death.

(1) seems perfectly appropriate. However, when we apply Yablo's conditions, we find that they are not satisfied. We can now see this in detail: first, let *hemlock* be the event that is essentially Socrates' drinking hemlock; and let *poison* be the event that is essentially Socrates' drinking a lethal poison and only accidentally his drinking hemlock. We now find that *hemlock* is not required for Socrates'

Footnote 9 continued
there may in principle be cases where the reverse does not hold. However, this difference between the two versions does not matter in any of the cases I consider here, or indeed in any other cases I can think of.

death: *poison* is a weakening of *hemlock*, and if *poison* had occurred without *hemlock*—that is, if Socrates had drunk a lethal poison without drinking hemlock—Socrates would still have died. On the natural assumption that "Socrates' drinking hemlock" refers to *hemlock*, the idea that proportionality is a necessary condition for causation therefore yields the result that (1) is false, since *hemlock* is not proportional to Socrates' death. As Bontly writes:

> "insisting that causes be in proportion to their effects has some rather surprising consequences. At first, we took it at face value that Socrates' drinking hemlock caused his death; now we find that even so straightforward a causal claim as that is, on Yablo's view, most likely false. For the proportionality principle compels us to reject any but the least specific event that is still enough for the effect in question […] In the process, it forces us to abandon a great many causal claims that, in many quite ordinary explanatory contexts, we would be happy to assert." (Bontly, 2005, p. 341)

This is already bad news. It gets even worse when we consider the fact that, as McDonnell (2017, pp. 1470–71) has noted, a proportionality requirement implausibly *removes* the problem of redundant causation. Consider a classic case of late preemption: [10] Suzy and Billy are both throwing rocks at a window. Suzy throws her rock a moment before Billy throws his, and her rock hits the window first. The window shatters, and once Billy's rock reaches the place where the window pane used to be, it meets only thin air. There is nearly unanimous agreement about what the correct verdict is in this case: Suzy's throw is a cause of the window-breaking; Billy's is not. It has been a longstanding challenge to find conditions for causation that can distinguish a genuine cause, such as Suzy's throw, from a preempted backup, such as Billy's throw. But the verdict itself has not been in question. As it turns out, however, the idea that proportionality is a necessary condition for causation yields the verdict that neither Suzy's throw nor Billy's throw is a cause of the window-shattering. How so? Well, among the events that occur in the spacetime region where Suzy and Billy are, let *Suzy's throw* be the event that is essentially Suzy's throwing her rock at the window, let *Billy's throw* be the event that is essentially Billy's throwing his rock at the window, and let *someone's throw* be the event that is essentially Suzy's or Billy's throwing a rock at the window. We now find that *Suzy's throw* is not required for the window-shattering: *someone's throw* is a weakening of *Suzy's throw*, and if *someone's throw* had occurred without *Suzy's throw*, the window would still have shattered. Similarly, *Billy's throw* is not required for the window-shattering. Instead, the idea that proportionality is a necessary condition for causation yields the verdict that *someone's throw* is the cause of the shattering—neither *Suzy's throw* nor *Billy's throw* is a cause, since neither is proportional. That result is unacceptable. If the proportionality constraint forces us to reject *Suzy's throw* as a cause of the shattering, then so much the worse for the proportionality constraint.

What the late preemption case illustrates is that a proportionality constraint may sometimes force us to go to a level of detail—e.g. the level of *someone's throw*, rather than the more specific level of *Suzy's throw* and *Billy's throw*—where the interesting

[10] McDonnell uses a case of early preemption to illustrate the same point.

causal structures we thought we were investigating simply drop from view. This does not merely happen in cases of redundant causation. As Bernstein (2014) points out, it also happens when we apply a proportionality constraint to the problem of profligate omissions (for such an application, see Dowe, 2010). Let me illustrate the problem with a standard case: Billy has promised to water Suzy's flowers while she is away on holiday. However, he fails to do so and the flowers die. The Queen of England also does not water the flowers. If either Billy or the Queen had watered the flowers, they would have blossomed. The intuitive verdict about the case seems clear: Billy's failure to water the flowers was a cause of their death; the Queen's failure to water them was not. A growing number of authors accept these verdicts at face value and take up the challenge of capturing what makes Billy's failure to water the flowers a cause, when the Queen's failure to do so is not.[11] If we take proportionality to be a necessary condition for causation, however, it turns out that both Billy's failure to water and the Queen's failure to water are at the wrong level of detail. To see this, let *Billy's failure* be the event that is essentially the flowers' receiving no water from Billy; let *Queen's failure* be the event that is essentially the flowers' receiving no water from the Queen; and let *no water* be the event that is essentially the flowers' receiving no water from anyone. We then find that *Billy's failure* is not enough for the flowers' death: *no water* is a strengthening of *Billy's failure*, and if *Billy's failure* had occurred without *no water*—that is, if Billy had failed to water the flowers, but someone else had watered them—then the flowers would not have died. The proportionality constraint thus yields the result that neither *Billy's failure* nor *Queen's failure* is a cause of the flowers' death—instead, the one proportional cause is simply *no water*. This wipes out what looked like an interesting causal structure, differentiating between Billy and the Queen, and instead gives us a simple causal relationship at a different level of detail, between *no water* and the death of the flowers. As in the case of late preemption, the proportionality constraint is not merely contradicting the commonsense judgement that Billy's failure to water the flowers caused their death; it is forcing us to abandon what looked like rich and interesting distinctions and giving us an uninformative truism instead.

4 Proportionality relative to a domain of events

The kind of objection to a proportionality constraint on causation that I have set out in the previous section is widespread in the literature: it is often repeated that a proportionality constraint would force us to reject perfectly good ordinary causes in favour of overly disjunctive and abstract events. As Franklin-Hall writes, the causes recommended by a proportionality constraint "are pitched at such great heights as to induce a kind of explanatory hypoxia, specifying far too little about what actually brought about the explanandum event to be very explanatory of it" (Franklin-Hall, 2016, p. 568). In later work, Yablo offers a response to this kind of objection by suggesting that "proportionality is not pursued at all costs but traded off against naturalness" (Yablo, 2003, p. 326; cf. Yablo, 2005, p. 466).

[11] See e.g. McGrath (2005) and Blanchard and Schaffer (2017).

Both this objection to proportionality and Yablo's later response ignore an alternative way to address the problem, which in fact was suggested by Yablo himself (1992a).[12] The fact that this suggestion has not been followed up in the literature so far, neither by Yablo nor anyone else, is puzzling. In the following, I show how Yablo's suggestion can be developed to yield a sophisticated account of proportionality, allowing us to reconcile a proportionality constraint with the truth of ordinary causal claims, such as "Socrates' drinking hemlock caused his death."

Yablo's suggestion begins with the observation that the definitions of being *required* and of being *enough* quantify over weakenings or strengthenings of an event x: x is required for y just in case *for all* $x^- < x$ …; x is enough for y just in case *for all* $x^+ > x$ … As the conditions are written, there is no further restriction on the domain of quantification—we are quantifying over absolutely *all* weakenings and strengthenings of x. However, the fact that both definitions use quantification opens up the possibility of *relativizing* proportionality to a domain of quantification. Yablo notes that "[proportionality's] demands intensify as its quantifiers range over more and more events" (Yablo, 1992a, p. 420) and suggests that the quantificational domain may vary depending on the context (Yablo, 1992a, pp. 424–5).

We can make this suggestion explicit as follows: for any domain of events D, we define being *required* and being *enough* relative to D, as follows:

> *Required*: x is *required* for y relative to domain D just in case for all x^- in D, such that $x^- < x$, if x^- had occurred without x, then y would not have occurred.
>
> *Enough*: x is *enough* for y relative to domain D just in case for all x^+ in D, such that $x^+ > x$, if x had occurred without x^+, then y would still have occurred.

On the basis of this, we get an account where proportionality is itself relativized to a domain of events:

> *Proportional*: x is *proportional* to y relative to domain D just in case
> (a) x is required for y relative to D, and
> (b) x is enough for y relative to D.

From here, there are two ways one might go. First, one might maintain that causation itself is not relative to a domain of events. To get a proportionality constraint on causation, one would then need to claim that there is some privileged domain of events such that x is a cause of y *only if* x is proportional to y relative to this privileged domain. Second, one might maintain that causation itself *is* relative to a domain of events. On this picture, a proportionality constraint would look like this: x is a cause of y relative to domain D *only if* x is proportional to y relative to domain D. In the following, I will argue that this second option is quite attractive. Let me begin by spelling out in more detail what it involves.

[12] The suggestion I am interested in here appears in the more detailed of Yablo's two proportionality papers from 1992, namely "Cause and essence" (1992a); it is not included in "Mental causation" (1992b).

We may think of a domain of events as an extra causal relatum.[13] Thus, the causal relation takes the form:

c is a cause of *e* relative to domain D

Correspondingly, a complete causal claim takes the form:

"*c* is a cause of *e* relative to domain D"

Usually, however, our causal claims take the binary form "*c* is a cause of *e*": we rely on context to supply the relevant domain of events. The relativization to a domain of events thus introduces an element of context-relativity to our causal claims. In the following section, I show how this allows us to reconcile a proportionality constraint with our commonsense verdicts about the trouble cases presented in Sect. 3.

5 Return to the cases

The first thing to note about the cases is that in each case, there is a domain D such that the purported cause is proportional to the effect relative to D.

In the case of Socrates' death, we have seen that when we consider *all* events, we find that *hemlock* is not proportional to Socrates' death because it is not required: if *poison* had occurred without *hemlock*, Socrates would still have died. However, there is a restricted domain relative to which *hemlock* is required: the restricted domain D, which does not contain any weakenings of *hemlock*, such as *poison*. Relative to D, it is trivially true that *hemlock* is required for Socrates' death, since there simply are no weakenings of *hemlock* in D. Thus, *hemlock* is proportional to Socrates' death relative to D.

In our late preemption case, we have similarly seen that when we consider *all* events, we find that *Suzy's throw* is not required for the window-shattering: if *someone's throw* had occurred without *Suzy's throw*, the window would still have shattered. However, when we consider a restricted domain D that simply leaves out all weakenings of *Suzy's throw*, we find that it is trivially true that *Suzy's throw* is required for the window-shattering. Thus, *Suzy's throw* is proportional to the window-shattering relative to D.

Finally, in our omission case, we have seen that when we consider *all* events, we find that *Billy's failure* is not enough for the death of the flowers: if *Billy's failure* had occurred without *no water*, the flowers would not have died. Here, *Billy's failure* does not satisfy the condition because there are strengthenings of *Billy's failure*, such as *no water*, that do better. However, when we consider a restricted domain D that leaves out all strengthenings of *Billy's failure*, it is trivially true that *Billy's failure* is enough

[13] It has frequently been suggested that causation may have further relata, in addition to the cause *c* and effect *e*. Within the causal modelling approach to causation (exemplified in the work of Halpern, Hitchcock, Pearl, Woodward, and others; see in particular Woodward, 2003), it is commonly suggested that causation is relative to a causal model (see e.g. Halpern & Pearl, 2005, p. 845). Furthermore, several authors have suggested that causation is contrastive and that the causal relation has contrasts to the cause and/or the effect as further relata (see e.g. Hitchcock, 1996a, 1996b; Maslen, 2004; Northcott, 2008; Schaffer, 2005).

for the death of the flowers. Thus, *Billy's failure* is proportional to the death of the flowers relative to D.

In all these cases, we can find a domain relative to which the purported cause is proportional to the effect, simply by leaving out the events—whether they are weakenings or strengthenings—that overturned that verdict when we were considering *all* events. This is not surprising: as I mentioned above, one may think of proportionality as a competition between events that are weakenings and strengthenings of each other. Introducing domain relativity means an event does not automatically enter the competition: before the competition even begins, there is the question of which events get to compete by being included in the relevant domain. By keeping some events out of the competition, we can ensure that any given event is the winner among those that get to compete: even if you are only the *n*th best in the world, you still win if all those who are better than you get excluded from the competition.

Of course, showing that there is a domain D relative to which e.g. *hemlock* is proportional to Socrates' death is not enough to capture our commonsense verdicts about the cases. To do so, we also need to show that D is the contextually relevant domain: the verdict we want is that it is appropriate to assert (1):

(1) Socrates' drinking hemlock caused his death.

On my proposal, (1) expresses the proposition (1*), where D is the contextually relevant domain of events:

(1*) Socrates' drinking hemlock caused his death relative to domain D.

In order to show that it is appropriate to assert (1), I therefore need to show that (1*) is true—and this requires showing that *hemlock* is proportional to Socrates' death relative to the contextually relevant domain.

Do we have any reason to think that the contextually relevant domain D is one that excludes all weakenings of *hemlock*, such that *hemlock* is proportional to Socrates' death relative to D? In fact, we do. The reason is to be found in the phenomenon of accommodation. As Lewis notes,

> "conversational score does tend to evolve in such a way as is required in order to make whatever occurs count as correct play." (Lewis, 1979, p. 347)

More carefully, we may think of a context as a conversational score. A conversational score may be understood as an *n*-tuple, where each of the *n* components keeps track of some aspect of the context—standards of precision, presuppositions, question under discussion, etc. When someone says something that requires some component of the conversational score to have a value within a certain range, in order for what is said to be acceptable, we typically accommodate by updating the conversational score so that the relevant component *does* take on a value in the required range. Suppose, for example, that Tom says, "My sister is picking me up at the train station." This presupposes that Tom *has* a sister: if the statement is to be appropriate, this presupposition needs to be in place. And as Lewis notes:

> "it's not as easy as you might think to say something that will be unacceptable for lack of required presuppositions. Say something that requires a missing presupposition, and straightway that presupposition springs into existence, making what you said acceptable after all." (Lewis, 1979, p. 339)

Even if the presupposition is not in place when Tom speaks, we typically accommodate by updating the conversational score on the fly: in this case, we add in the presuppositions that are needed in order to make what is said acceptable.

This brings me back to the hemlock case. On my proposal, one of the components of the conversational score is the relevant domain of events. If someone makes a causal claim that requires a domain within a certain range in order to be appropriate, I suggest that we typically accommodate by updating the conversational score so that the relevant domain does indeed fall within the required range. Consider (1), for example:

(1) Socrates' drinking hemlock caused his death.

On my proposal, (1) makes a true claim only if the contextually relevant domain D excludes all weakenings of *hemlock*. For this very reason, a hearer will typically accommodate by updating the conversational score so that the relevant domain D *does* exclude all weakenings of *hemlock*.[14] Once we have established, as we did above, that there *are* domains relative to which (1) is true, accommodation typically ensures that one of these domains is indeed the contextually relevant one.

In some cases, this accommodation starts from a "clean slate," where the component of the conversational score that keeps track of the relevant domain of events has not yet received any value. In other cases, the conversational score may already have some particular value for the domain of events at the time of the utterance, and we may then accommodate by changing this value on the fly. This may be illustrated by a variation of our preemption case. Suppose we are discussing whether Suzy or Billy broke the window, and we have one key piece of evidence: whoever broke the window was throwing with their right hand. Suppose further that Suzy is right-handed and threw

[14] I am here assuming that the hearer already knows the relevant facts about the case, in particular that Socrates would still have died if he had ingested some other lethal poison. Based on this, it will be clear to the hearer that (1) only makes a true claim relative to a domain D that excludes all weakenings of *hemlock*, and that accommodation therefore requires updating the conversational score so that the contextually relevant domain D *does* exclude all weakenings of *hemlock*.What if the hearer is not aware of the relevant facts about the case? If so, it may be an open possibility to her that (1) might also express a true claim relative to a larger domain D* which includes *poison*: for all she knows, the speaker might be making the startling claim that *hemlock* was required for Socrates' death relative to the larger domain D*. Maybe Socrates had built up a tolerance of every other lethal poison at the Athenians' disposal; thus, if he had ingested a lethal poison (i.e. lethal to other people) without drinking hemlock, he would not have died. When a hearer is in this position, she does not know what (if anything) is needed to get the utterance to come out true; therefore, she does not know what accommodation requires. If she cannot make an informed guess, she may need to ask for clarification—or risk misunderstanding what was said. This is indeed a general feature of accommodation. To take a different example, consider the utterance that "France is hexagonal" (cf. Lewis, 1979, p. 352; Ball & Huvenes, 2021). Given the actual shape of France, this utterance makes a true claim only if the contextually relevant standards of precision are quite low. A hearer who knows enough about the shape of France—in particular, that its coastline has many bays and outcrops—will immediately realize this. However, to a hearer who is entirely unfamiliar with the geography of France (or indeed, geography in general), it may be an open possibility that the speaker is making the startling claim that France is hexagonal relative to a much more demanding standard of precision. In some cases, this may lead to misunderstandings. I am grateful to an anonymous reviewer for prompting me to clarify this.

with her right hand, while Billy is left-handed and threw with his left hand. If we have just been discussing this, it is highly plausible that the relevant domain of events D includes all of the following events: *Suzy's throw*, *Suzy's right-handed throw*, *Billy's throw*, and *Billy's left-handed throw*, where *Suzy's right-handed throw* is the event that is essentially Suzy's throwing with her right hand, and correspondingly in the case of *Billy's left-handed throw*. Suppose that you then say, "Suzy's throw caused the window to shatter." Relative to D, this claim is false: supposing that Suzy could not hit the broad side of a barn with her left hand, *Suzy's right-handed throw* defeats *Suzy's throw* in a proportionality contest. The reason: *Suzy's throw* is not enough for the window-shattering—if *Suzy's throw* had occurred without *Suzy's right-handed throw*, the window would not have shattered (since Suzy would have missed). For this very reason, the hearer will typically accommodate by updating the contextually relevant domain from D to a smaller domain D*, which does not include *Suzy's right-handed throw*. With D* as the contextually relevant domain, your utterance now comes out true.[15]

In this way, my proposal can agree with our judgements about ordinary causal claims, as exemplified above, while maintaining that proportionality is a necessary condition for causation: in each case, there exist domains relative to which the purported cause is proportional to the effect; and accommodation ensures that one of these domains is the contextually selected one.

At this point, however, you might have the opposite worry: not that a proportionality constraint cannot make room for our ordinary causal claims, but that it can make room for *any* causal claim. Given that we can always restrict which events enter the competition, so that we get the winner we want, does a relativized proportionality constraint add anything—that is, does it do any work? Or is the price of relativizing that the proportionality constraint has lost its bite? Could we do just as well without it? In the following section, I take up this challenge and show that a relativized proportionality constraint can still do important work.

6 The work that a relativized proportionality constraint can do

In this section, I show that a relativized proportionality constraint can still offer elegant solutions to the problems that might motivate us to adopt a proportionality constraint in the first place: it allows us to deny rampant overdetermination, and it handles cases where we need to choose between candidate causes at different levels of detail. In addition, a relativized proportionality constraint can explain phenomena of retraction and negotiation that a non-relativized proportionality constraint struggles to capture.

Let us begin with overdetermination. According to common sense, overdetermination is a rare phenomenon. It only occurs when things line up just right—when two rocks hit a window, or two bullets hit a victim, at exactly the same time. However, without some kind of proportionality constraint, we seem forced to say that overdetermination is not rare at all—in fact, there is overdetermination all the time.[16] According

[15] I am grateful to an anonymous reviewer for prompting me to clarify this.

[16] For an overview of the debate, see e.g. Paul and Hall (2013), pp. 155–161.

to common sense, for example, Socrates' death was not overdetermined. But without a proportionality constraint, we seem forced to say it was: it was caused by Socrates' drinking hemlock, and it was caused by Socrates' drinking a lethal poison, and both of these causes were sufficient (in the circumstances) to bring it about.

Worries about rampant overdetermination are one of the main motivations for adopting a proportionality constraint on causation. A relativized proportionality constraint provides an elegant solution: it is true that *hemlock* is proportional to Socrates' death relative to one domain D (where D does not include any weakenings of *hemlock*), and it is true that *poison* is proportional to Socrates' death relative to a larger domain D* (which includes both *poison* and *hemlock*). However, it is not true relative to *any* domain that *hemlock* and *poison* are *both* proportional to Socrates' death. Thus, it is not true relative to *any* domain that Socrates' death was overdetermined. Relativized proportionality allows us to have our cake and eat it too: we can capture the commonsense verdict that it is acceptable to utter both (1) and (2), as long as they are uttered in different contexts so that accommodation can do its work, while *also* capturing the commonsense verdict that (1 + 2) is unacceptable, no matter which context it is uttered in:

(1) Socrates' drinking hemlock caused his death.
(2) Socrates' drinking a lethal poison caused his death.
(1 + 2) Socrates' drinking hemlock and Socrates' drinking a lethal poison caused his death.

This shows that a relativized proportionality constraint can do just as well as a non-relativized proportionality constraint when it comes to dealing with worries about rampant overdetermination.[17] Relativizing the proportionality constraint has not taken away its bite.

Moving on, consider questions such as the following: "was the object's being scarlet, the object's being red, or the object's being coloured a cause of Sophie's peck?" The intuitively correct answer is that "the object's being red was a cause of Sophie's peck." A proportionality constraint allows us to give this answer. Without a proportionality constraint, this kind of answer may not be available—instead, we might have to say that *scarlet*, *red*, and *coloured* are all causes (bringing us back to the problem of rampant overdetermination discussed above). This is a further motivation for adopting a proportionality constraint on causation. Once again, if this is your motivation for adopting a proportionality constraint, a relativized proportionality constraint does just as well as a non-relativized one: when we are asking "was the object's being scarlet, the object's being red, or the object's being coloured a cause of Sophie's peck?", the contextually relevant domain D clearly includes all three events, *scarlet*, *red*, and *coloured*. All three events are allowed to compete, so to speak, and the conditions for

[17] There are, of course, other ways to deny rampant overdetermination *without* accepting a proportionality constraint on causation. One might, for example, reject Yablo's extremely fine-grained event ontology and hold instead that e.g. "Socrates' drinking hemlock" and "Socrates drinking a lethal poison" are merely different ways to describe the very same event. My aim in this paper is not to argue that we *should* accept a proportionality constraint on causation in order to deal with overdetermination worries. Rather, I merely aim to show that *if* overdetermination worries motivate you to accept a proportionality constraint, a relativized proportionality constraint can do the job just as well as a non-relativized one. I am grateful to an anonymous referee for helping me to clarify my argument.

proportionality relative to D therefore do real work in picking out the winner—in this case, *red*.

This also means that the relativized proportionality constraint can do the work that Yablo wanted a proportionality constraint to do: when we ask whether Jones's mental state (being in pain) or his microphysical brain state caused him to flinch, the relativized proportionality constraint delivers the result that Jones's mental state was the cause; not his microphysical brain state. To see this, we need to first specify the two events: *mental state* is essentially Jones's being in a particular mental state, viz. being in pain; *brain state* is essentially Jones's being in a particular microphysical brain state. Yablo suggests that *mental state* is a weakening of *brain state*: just like the object's being red can be realized in a multitude of ways, one of which is the object's being scarlet, Jones's being in pain can be realized in a multitude of ways, one of which is his having this particular microphysical brain state (see e.g. Yablo, 1992a, p. 430). If we accept this, we find that *brain state* is not proportional to Jones's flinching, since *mental state* is a weakening of *brain state*, and if *mental state* had occurred without *brain state*—that is, if Jones had been in pain without being in this particular brain state—he would still have flinched. On the other hand, *mental state* is proportional to Jones's flinching. That was precisely the result Yablo wanted—namely, "a story in which a mental event emerges as *better* qualified than its physical basis for the role of cause" (Yablo, 1992a, p. 436).

Finally, a relativized proportionality constraint offers us a way to answer the following challenge: suppose that a platform will collapse if a weight of more than 1000 kg is placed on it. In fact, a weight of 1600 kg is placed on it. Letting $\text{Statement}_{1600\text{kg}}$ be the statement that "the weight's being 1600 kg caused the collapse," and $\text{Statement}_{>1000\text{ kg}}$ be the statement that "the weight's being more than 1000 kg caused the collapse," Maslen notes that

> "[Yablo] owes us an explanation of why initially it seems perfectly acceptable to say that the platform's weighing 1600 kg caused the collapse ($\text{Statement}_{1600\text{kg}}$) and we only tend to feel dissatisfied with it when we are offered an improved alternative ($\text{Statement}_{>1000\text{kg}}$)." (Maslen, 2017, p. 65)

This phenomenon is indeed difficult to explain on the basis of a non-relativized proportionality constraint: A non-relativized proportionality constraint yields the verdict that $\text{Statement}_{1600\text{kg}}$ is false, no matter whether we are considering $\text{Statement}_{>1000\text{ kg}}$ as an alternative or not. On a non-relativized account, $\text{Statement}_{1600\text{kg}}$ is false *tout court*, because the weight's being 1600 kg is not proportional to the platform's collapse (if the weight had been more than 1000 kg without being 1600 kg, the platform would still have collapsed). The phenomenon is similarly difficult to explain if we have no proportionality constraint at all: if $\text{Statement}_{1600\text{kg}}$ and $\text{Statement}_{>1000\text{ kg}}$ are both true, why would we be dissatisfied with $\text{Statement}_{1600\text{kg}}$ when we are offered $\text{Statement}_{>1000\text{ kg}}$ as an alternative? In what sense is $\text{Statement}_{>1000\text{ kg}}$ an improvement?

This is of course not to say that our reactions *cannot* be explained on the basis of a non-relativized proportionality constraint, or no proportionality constraint at all (see e.g. Maslen, 2017). However, it is worth noting that a relativized proportionality constraint offers a natural and straightforward explanation: let *1600 kg* be the event that

is essentially the weight's being 1600 kg, and let > *1000 kg* be the event that is essentially the weight's being more than 1000 kg. When someone utters $\text{Statement}_{1600\text{kg}}$, we will typically accommodate and update the conversational score so that the contextually relevant domain is a domain D which includes *1600 kg* but does not include any weakenings of *1600 kg*, such as > *1000 kg*. Relative to such a domain D, *1600 kg* is proportional to the platform's collapse, and thus $\text{Statement}_{1600\text{kg}}$ is true. However, once someone utters the alternative $\text{Statement}_{>1000\text{ kg}}$, it is natural to update the conversational score so that the contextually relevant domain is now a larger domain D* that includes both *1600 kg* and > *1000 kg*. Relative to D*, *1600 kg* is not proportional to the platform's collapse, since it is not required: if > *1000 kg* had occurred without *1600 kg*, the platform would still have collapsed. By contrast, > *1000 kg* is proportional to the platform's collapse relative to D*. Relative to the new contextually relevant domain D*, $\text{Statement}_{1600\text{kg}}$ is thus false, while $\text{Statement}_{>1000\text{ kg}}$ is true. This explains why we become dissatisfied with $\text{Statement}_{1600\text{kg}}$ when we are presented with $\text{Statement}_{>1000\text{ kg}}$.[18]

These examples show that a relativized proportionality constraint can still do important work: the relativization does not mean that the proportionality constraint loses its bite.

7 Choosing a domain

On my proposal, a given causal situation may be truly described in many different ways. Consider, for example, the case of Sophie the pigeon. Let D be a domain that includes the event *scarlet*, but does not include any weakenings of *scarlet*, such as *red* or *coloured*. Let D* be a larger domain that includes both *scarlet*, *red*, and *coloured*. Then both of the following causal relations hold, and the corresponding causal claims are therefore true: *scarlet* is a cause of Sophie's peck relative to D, and *red* is a cause of Sophie's peck relative to D* (since *scarlet* is proportional to Sophie's peck relative to D, *red* is proportional to Sophie's peck relative to D*, and I assume that other conditions for causation are satisfied).

So far, I have focused on the hearer: once an utterance has been made, the hearer typically accommodates by updating the conversational score as needed to make the utterance come out true. If someone says that "the object's being scarlet caused Sophie to peck," the hearer typically accommodates by updating the conversational score so that D is the contextually relevant domain; if someone says that "the object's being red caused Sophie to peck," the hearer typically accommodates by updating so that D* is now the contextually relevant domain. As a speaker, however, you face a choice: Which of these utterances should you make? Which of the many causal relations should you pay attention to and talk about? In the following, I suggest that the purpose of our causal inquiry plays a crucial role in answering these questions: a particular purpose—such as attributing responsibility or preventing future accidents—may be

[18] At the end of Sect. 7, I note that we may sometimes refuse to accommodate if one domain of events serves the purpose of our causal inquiry better than another. If we have already been presented with $\text{Statement}_{>1000\text{ kg}}$, and domain D* better serves the purpose of our causal inquiry, we might for example refuse to accommodate as needed to make $\text{Statement}_{1600\text{kg}}$ appropriate.

better served by one choice of domain than another. As a speaker, you should, other things being equal, choose a domain within the range of domains that best serve the purpose of the causal inquiry. Here are some examples:

Suppose first that the purpose of our causal inquiry is to attribute responsibility. To fulfil this purpose, we need to understand how individual people are causally related to the effect. This aim is best served by a domain that includes events that essentially involve named individuals—Suzy, Billy, the Queen, etc.—and which does not include competing strengthenings or weakenings of those events that do not essentially involve named individuals.

Consider, for example, the case of the window-shattering: let D be a domain that includes *Suzy's throw* and *Billy's throw*, but which does not include events such as *someone's throw*, since *someone's throw* does not involve a single, named individual. Let D* be a larger domain that includes both *Suzy's throw*, *Billy's throw*, and *someone's throw*. If the purpose of our causal inquiry is to attribute responsibility for the window-shattering, it is now clear that we can better fulfil this purpose by considering domain D: relative to domain D, we find that *Suzy's throw* is a cause of the shattering (since *Suzy's throw* is proportional to the window-shattering relative to D), while *Billy's throw* is not a cause of the shattering (since *Billy's throw* fails to satisfy other conditions of causation, given that it is a preempted backup). That is precisely the kind of answer that helps us to hold someone responsible—in this case Suzy. By contrast, we cannot fulfil our purpose by considering domain D*: here, we merely get the result that *someone's throw* was a cause of the window-shattering (since only *someone's throw* is proportional to the shattering relative to D*). That result is not at all helpful for holding some particular individual—Suzy, or Billy, or some third person—responsible for the shattering.

Similarly in the case of profligate omissions: let D be a domain that includes *Billy's failure* and *Queens's failure*, but leaves out events such as *no water*, which do not involve a single, named individual. Let D* be a larger domain that includes both *Billy's failure*, *Queen's failure*, and *no water*. Finally, assume that, setting proportionality constraints aside, *no water* and *Billy's failure* count as causes of the flowers' death, while *Queen's failure* does not. When our purpose is to attribute responsibility for the death of the flowers, it is now clear that we do better by considering domain D than by considering domain D*: relative to domain D, *Billy's failure* is proportional to the flowers' death. Looking at domain D therefore helps us to attribute responsibility: relative to D, *Billy's failure* is a cause of the flowers' death, while *Queen's failure* is not. Relative to domain D*, on the other hand, we find that only *no water* is proportional to the flowers' death. This does not at all help us to attribute responsibility to individuals.

When the purpose of our causal inquiry is to attribute responsibility, we therefore find that we have reason to prefer some domains over others. Other purposes may similarly be better served by some domains rather than others. To see how, let us return once again to the case of Sophie the pigeon. As we have seen, *scarlet* is a proportional cause of Sophie's peck relative to a domain D that includes *scarlet* but does not include any weakenings of *scarlet*, such as *red* or *coloured*. By contrast, *red* is a proportional cause of Sophie's peck relative to the larger domain D*, which includes both *scarlet*, *red*, and *coloured*. Yablo suggests that different purposes may

prompt us to either widen our domain, as we do when we move from D to D*, or contract it, as we do when we move from D* to D:

> "[Enlarging the domain] turns up the commensuration pressure on would-be causes. Relatively incidental features of the causal scene, distinctive though they might be of the actual progress of events, are worn away to reveal the steadier causal currents beneath. Such a strategy can of course be taken too far […] Practised in moderation, though, it brings on an agreeable broadening and deepening of causal judgement, what I described by saying that these judgements become less world- and more effect-driven. Sometimes, it is true, we are willing to accept a shallower causal story in return for more discriminating information about what took place; in that case an easing of commensuration pressures is called for and hence a reduced causal ontology [i.e. a smaller domain[19]]." (Yablo, 1992a, pp. 424–425)

On this picture, we are balancing two competing considerations when we are choosing a domain: enlarging the domain will allow more events into the proportionality competition, and the winner will therefore reveal more about what was *needed* for the effect—it will be closer to the ideal of including just the right amount of detail, neither too much nor too little. On the other hand, contracting the domain, so that it only includes events with more demanding conditions of occurrence (such as *scarlet*), reduces the extent to which the winner needs to live up to the ideal of including just the right amount of detail. What we gain instead is more information about what actually happened—as Yablo puts it, we learn more about how the effect's needs "were in fact met" (Yablo, 1992a, p. 424).

Depending on the purpose of our causal inquiry, we may weigh these considerations differently—we may take it to be more or less important to find out the details of what actually happened, compared to finding out what was *needed* for the effect to occur. When an engineer is considering the collapse of the platform, for example, the primary purpose of his causal inquiry may be to ensure that this kind of accident does not happen again. To achieve this, he should emphasize finding out what was *needed* for the effect to occur, never mind how exactly that need was met. His purpose is therefore served better by considering the larger domain D*, relative to which we find that $> 1000\ kg$ is a cause of the collapse: focusing on this is more helpful for avoiding future accidents than focusing on the smaller domain D, relative to which we find that $1600\ kg$ is a cause of the collapse. By contrast, a historian investigating Socrates' death might place relatively more weight on how it actually happened, rather than merely what was needed for it to occur. Given his purpose, it might therefore be preferable to consider the smaller domain D, relative to which *hemlock* is a cause of Socrates' death, rather than the larger domain D*, relative to which *poison* is a cause of Socrates' death. In Yablo's terms, what the historian is after may well be a more world-driven cause—something he gets by contracting the domain.

This understanding of the considerations that go into a speaker's choice of domain may also help us to get a more nuanced understanding of accommodation. I said above that when someone makes a causal claim, we typically accommodate by updating

[19] See Yablo (1992a), p. 420: "quantificational domain— what I will call *causal ontology*."

the conversational score, so that the causal claim is true relative to the (updated) contextually relevant domain. In some cases, however, we may refuse to accommodate. Instead, we may counter with a causal claim of our own. If one engineer says that "the weight's being 1600 kg caused the platform to collapse," a second engineer might refuse to accommodate and instead counter by saying "the weight's being more than 1000 kg caused the platform to collapse." If one historian says that "Socrates' drinking poison caused his death," another historian might refuse to accommodate and counter by saying "Socrates' drinking hemlock caused his death."[20] If a teacher says "someone's throwing a rock caused the window to shatter," a second teacher might counter by saying "Suzy's throwing a rock caused the window to shatter." After that, a negotiation might ensue, ostensibly about what caused the collapse or Socrates' death or the window-shattering but in reality about what the contextually relevant domain should be.

We are now in a position to better understand why someone might refuse to accommodate, and what considerations might go into the ensuing negotiation about the choice of domain: since different domains are more or less well suited to the different purposes of our causal inquiries, you may refuse to accommodate if you feel that your interlocutor's chosen domain is ill-suited to the purpose of your inquiry. The ensuing negotiation may in turn be about determining which domain will in fact best serve the purpose.

8 Conclusion

One of the main objections to a proportionality constraint on causation is that many ordinary and seemingly perfectly acceptable causal claims cite causes that are not proportional to their effects. In this paper, I have argued that a proportionality constraint may be reconciled with the truth of our ordinary causal claims by developing a suggestion that is already present in Yablo's early work on proportionality (1992a): that proportionality, and causation itself, is relative to a domain of events.

At the metaphysical level, this proposal amounts to the claim that the causal relation has an extra relatum—namely, a domain of events. At the level of language, the proposal introduces a new way in which causal claims are sensitive to context: when the contextually relevant domain of events is D, "*c* is a cause of *e*" expresses the complete causal claim that *c* is a cause of *e* relative to D; when the contextually relevant domain of events is D*, "*c* is a cause of *e*" expresses the complete causal claim that *c* is a cause of *e* relative to D*.

If you are at all tempted by a proportionality constraint on causation, I think you should adopt this relativized notion of proportionality. Its costs are lower: while a standard proportionality constraint rejects our ordinary causal claims as false, a relativized proportionality constraint accommodates the verdict that they are true. And its benefits are greater: a relativized proportionality constraint does as well as a standard proportionality constraint in dealing with rampant overdetermination and choices

[20] As Bontly (2005, pp. 341–42) notes, a coroner's report would similarly require precision—even to the point of specifying the species of hemlock (*Conium maculatum*, as opposed to *Cicuta douglasii*).

between weakenings or strengthenings of the same event (as when we ask whether *scarlet* or *red* was a cause of Sophie's peck). In addition, a relativized proportionality constraint can explain phenomena that a standard proportionality constraint cannot: it can explain why we are initially satisfied with a claim such as "the weight's being 1600 kg caused the platform's collapse," and only get dissatisfied with it once we are presented with an alternative, such as "the weight's being more than 1000 kg caused the collapse."[21]

Funding Open access funding provided by Umea University.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Ball, D., & Huvenes, T. (2021). A puzzle about accommodation and truth. *Philosophical Studies*. https://doi.org/10.1007/s11098-021-01693-2

Bernstein, S. (2014). Two problems for proportionality about omissions. *Dialectica, 68*, 429–441. https://doi.org/10.1111/1746-8361.12071

Blanchard, T., & Schaffer, J. (2017). Cause without default. In H. Beebee, C. Hitchcock, & H. Price (Eds.), *Making a difference: Essays on the philosophy of causation* (pp. 175–214). Oxford University Press.

Bontly, T. D. (2005). Proportionality, causation, and exclusion. *Philosophia, 32*, 331–348. https://doi.org/10.1007/BF02641629

Dowe, P. (2010). Proportionality and omissions. *Analysis, 70*, 446–451. https://doi.org/10.1093/analys/anq033

Franklin-Hall, L. R. (2016). High-level explanation and the interventionist's 'variables problem.' *British Journal for the Philosophy of Science, 67*, 553–577. https://doi.org/10.1093/bjps/axu040

Halpern, J. Y., & Pearl, J. (2005). Causes and explanations: A structural-model approach. Part I: Causes. *British Journal for the Philosophy of Science, 56*, 843–887. https://doi.org/10.1093/bjps/axi147

Hitchcock, C. (1996a). Farewell to binary causation. *Canadian Journal of Philosophy, 26*, 267–282. https://doi.org/10.1080/00455091.1996.10717454

Hitchcock, C. (1996b). The role of contrast in causal and explanatory claims. *Synthese, 107*, 395–419. https://doi.org/10.1007/BF00413843

Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic, 8*, 339–359. https://doi.org/10.1007/BF00258436

Lewis, D. (1986). Events. In D. Lewis, *Philosophical papers, volume II*. Oxford University Press.

List, C., & Menzies, P. (2009). Nonreductive physicalism and the limits of the exclusion principle. *Journal of Philosophy, 106*, 475–502. https://doi.org/10.5840/jphil2009106936

[21] I am grateful to Derek Ball, Katherine Hawley, Ned Hall, and L. A. Paul for fruitful comments and discussions during the early stages of developing this idea. I would like to thank Bram Vaassen and Alexander Sandgren for stimulating discussions of Yablo's treatment of proportionality and Torfinn Huvenes for organizing a reading group on accommodation at just the time when I needed to think more about it. Finally, I am grateful to an anonymous referee for this journal for many helpful comments.

Maslen, C. (2004). Causes, contrasts, and the nontransitivity of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 341–357). MIT Press.
Maslen, C. (2017). Pragmatic explanations of the proportionality constraint on causation. In H. Beebee, C. Hitchcock, & H. Price (Eds.), *Making a difference: Essays on the philosophy of causation* (pp. 58–72). Oxford University Press.
McDonnell, N. (2017). Causal exclusion and the limits of proportionality. *Philosophical Studies, 174*, 1459–1474. https://doi.org/10.1007/s11098-016-0767-3
McDonnell, N. (2018). Transitivity and proportionality in causation. *Synthese, 195*, 1211–1229. https://doi.org/10.1007/s11229-016-1263-1
McGrath, M. (1998). Proportionality and mental causation: A fit? *Noûs, 32*, 167–176. https://doi.org/10.1111/0029-4624.32.s12.7
McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies, 123*, 125–148. https://doi.org/10.1007/s11098-004-5216-z
Northcott, R. (2008). Causation and contrast classes. *Philosophical Studies, 139*, 111–123. https://doi.org/10.1007/s11098-007-9105-0
Paul, L. A., & Hall, N. (2013). *Causation: A user's guide*. Oxford University Press.
Sartorio, C. (2010). The Prince of Wales problem for counterfactual theories of causation. In A. Hazlett (Ed.), *New waves in metaphysics* (pp. 259–276). Palgrave Macmillan.
Schaffer, J. (2005). Contrastive causation. *Philosophical Review, 114*, 327–358. https://doi.org/10.1215/00318108-114-3-327
Shapiro, L. A., & Sober, E. (2012). Against proportionality. *Analysis, 72*, 89–93. https://doi.org/10.1093/analys/anr135
Weslake, B. (2013). Proportionality, contrast and explanation. *Australasian Journal of Philosophy, 91*, 785–797. https://doi.org/10.1080/00048402.2013.788045
Weslake, B. (2017). Difference-making, closure, and exclusion. In H. Beebee, C. Hitchcock, & H. Price (Eds.), *Making a difference: Essays on the philosophy of causation* (pp. 215–231). Oxford University Press.
Woodward, J. (2003). *Making things happen: A theory of causal explanation.* Oxford University Press.
Woodward, J. (2015). Interventionism and causal exclusion. *Philosophy and Phenomenological Research, 91*, 303–347. https://doi.org/10.1111/phpr.12095
Woodward, J. (2021). Explanatory autonomy: The role of proportionality, stability, and conditional irrelevance. *Synthese, 198*, 237–265. https://doi.org/10.1007/s11229-018-01998-6
Yablo, S. (1992a). Cause and essence. *Synthese, 93*, 403–449. https://doi.org/10.1007/BF01089276
Yablo, S. (1992b). Mental causation. *Philosophical Review, 101*, 245–280. https://doi.org/10.2307/2185535
Yablo, S. (2003). Causal relevance. *Philosophical Issues, 13*, 316–328. https://doi.org/10.1111/1533-6077.00016
Yablo, S. (2005). Prime causation. *Philosophy and Phenomenological Research, 70*, 459–467. https://doi.org/10.1111/j.1933-1592.2005.tb00541.x

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