

Halfway Proportionality

Bram Vaassen

May 10, 2022

Penultimate draft, please cite published version

Forthcoming in *Philosophical Studies*: 10.1007/s11098-022-01803-8

Abstract

According to the so-called ‘proportionality principle’, causes should be proportional to their effects: they should be both enough and not too much for the occurrence of their effects. This principle is the subject of an ongoing debate. On the one hand, many maintain that it is required to address the problem of causal exclusion and take it to capture a crucial aspect of causation. On the other hand, many object that it renders accounts of causation implausibly restrictive and often reject the principle wholesale. I argue that there is exaggeration on both sides. While one half of the principle is overly demanding, the other half is unobjectionable. And while the unobjectionable half does not block exclusion arguments on its own, it provides a nuanced picture of higher-level causation, fits with recent developments in philosophy of causation, and motivates adjustments to standard difference-making accounts of causation. I conclude that at least half of the proportionality principle is worth taking seriously.

Keywords: proportionality; causation; causal explanation; exclusion arguments; higher-level causation; difference-making; interventionism; contrastivism

1 Introduction

According to the so-called ‘proportionality principle’, causes should be proportional to their effects: they should be both enough and not too much for the occurrence of their effects. This principle divides philosophers writing on causation. Its proponents, including Yablo (1992a,b, 1997, 2003); Shoemaker (2001); List and Menzies (2009); Raatikainen (2010), and Zhong (2014, 2020a,b), take it to capture a crucial aspect of causation and maintain that it is required to address the problem of causal exclusion, i.e. the problem that leptons and quarks threaten to causally exclude familiar causes like colds and hurricanes. Opponents, including McGrath (1998); Bontly (2005); Shapiro and Sober (2012); Weslake (2013); Hoffmann-Kolss (2014); Maslen (2017); McDonnell (2017); Baumgartner (2018) and Woodward (2021), object that the proportionality principle is too demanding and often reject the principle wholesale. They typically accept that proportional causes tend to provide better explanations than non-proportional causes (Weslake, 2013; McDonnell, 2017; Woodward, 2021), or that citing non-proportional causes tends to violate conversational maxims (Bontly, 2005), but they deny that violations of the proportionality principle provide a ground for not being a cause.

I argue that there is exaggeration on both sides. While one half of the principle is overly demanding, the other half is unobjectionable. First, I introduce the proportionality principle (Section 2). I then discuss two standard objections and argue that one half of the proportionality principle can withstand them (Section 3). In Section 4, I discuss how this half is supported by recent work on the importance of (in)sensitivity for causation and discuss two consequences of adopting this ‘halfway’ proportionality principle: it provides a nuanced and interesting picture of higher-level causation, despite having limited force against the problem of causal exclusion, and it motivates adjustments to the difference-making accounts of causation proposed by Lewis (1973b, 1979), Schaffer (2005), and

Woodward (2003, 2021). I conclude that the halfway proportionality principle is worth taking seriously.

2 Proportionality

We can start our discussion with a familiar observation: some phenomena impose strictly stronger requirements on worlds than others. For example, an object being scarlet imposes a strictly stronger requirement on the world than its being red, because there are worlds where that object is red, but not scarlet, and there are no worlds in which the object is scarlet, but not red. Consequently, knowing that an object is scarlet provides me with strictly more information about the actual world than knowing that it is red, as knowledge of its scarletness allows me to exclude possible worlds that are not excluded by knowledge of its redness, but the reverse does not hold. Let us say that phenomena standing in such ‘a strictly stronger than’ relation are ‘intimately’ related: two phenomena are intimately related if and only if one of them is strictly stronger than the other.

Phenomena that are intimately related appear to be distinct in at least two ways. First, they have different identity conditions. An object’s being red is not the same as its being scarlet, as there are many ways to be red without being scarlet, such as being crimson or being auburn (cf. Pereboom, 2002; Paul, 2006, 2007).¹ Second, intimately related phenomena can appear to occupy different causal roles. My sweater’s being scarlet might cause you to like it, without its being red causing you to like it. This would be the case if you exclusively like scarlet sweaters, and thus would have disliked my sweater if it had been any

¹As noted by an anonymous referee, one might worry that accepting the distinctness of intimately related phenomena leads to an overpopulated ontology. My sweater’s being red is one phenomenon, my sweater’s being scarlet another, my sweater’s being this exact shade of scarlet another still, and so on. In response, proponents of such an abundant ontology argue that there is no significant cost in multiplying non-fundamental phenomena, such as being red and being scarlet, as long as these are built out of a sparse base of fundamental phenomena, such as leptons and quarks having spin and location (e.g. Cameron, 2010; Schaffer, 2015).

other shade of red. The proportionality principle provides a systematic approach to the apparent causal differences between intimately related phenomena.

According to the proportionality principle, intimately related phenomena automatically enter into a cut-throat competition for causal status. That is to say, if a phenomenon P causes a given effect E , then there is no phenomenon $P+$ that is strictly stronger than P and causes E and there is no phenomenon $P-$ that is strictly weaker than P and causes E .² To use a popular example, if an object being red causes Sophie, a pigeon trained to peck exclusively at red objects, to peck, then the proportionality principle states that neither its being scarlet nor its being coloured can be a cause of her pecking (cf. Yablo, 1992b, p. 257). More precisely, the proportionality principle dictates that the outcome of this causal competition is decided as follows:

Proportionality P causes E only if

- (i) For any $P+$, strictly stronger than P , P screens off E from $P+$
- (ii) For any $P-$, strictly weaker than P , $P-$ does not screen off E from P

Where the screening off relation is understood as follows:

Screening Off For any three non-identical phenomena P , Q and E , P screens off E from Q if and only if, E would still have occurred if P had occurred without Q .

For example, the redness of an object screens off Sophie's pecking from the object's being scarlet, if and only if she would have pecked at the object in nearby worlds in which it is red, but not scarlet.

PROPORTIONALITY captures the idea that causes should be enough, but not too much for their effects. Being scarlet is too much to cause Sophie's pecking;

²Where $P-$ is strictly weaker than P iff P is strictly stronger than $P-$.

it contains more than is relevant. This is evident from the fact that the redness screens off the pecking from the scarletness: if the pebble had been red but not scarlet, Sophie would still have pecked. Being coloured is not enough. This is evident from the fact that the object's being coloured fails to screen off the pecking from its being red: if the pebble had been coloured but not red, Sophie would not have pecked. Redness gets it just right, it screens off strictly stronger competitors and is not screened off by its strictly weaker competitors.³

Similar examples abound in the proportionality literature.⁴ Consider Zsa Zsa Gabor speeding past the speedometer and subsequently getting a ticket (Yablo, 1992b, p. 275). She got a ticket because she *sped* past, not because she *drove* past. Or consider Socrates drinking hemlock and subsequently dying. He died because he *drank* hemlock, his *guzzling* it did not contribute (Yablo, 1992b, p. 278). Consequently, we get the following claims:

- (1) Speeding past caused Zsa Zsa Gabor to get a ticket.
- (2) Driving past did not cause Zsa Zsa Gabor to get a ticket.

and

- (3) Guzzling hemlock did not cause Socrates' death.
- (4) Drinking hemlock caused Socrates' death.

The proportionality principle aligns with these judgments. As prescribed by (i), the strictly weaker phenomenon *driving by* fails to screen off speeding tickets from *speeding by* and, as prescribed by (ii), the strictly weaker phenomenon of Socrates drinking hemlock screens off his death from his guzzling it. At first glance, the results all check out.

Proponents of PROPORTIONALITY maintain that we should take these results at face value. Opponents maintain that the mutual acceptability of pairs like

³Though see Section 3.

⁴See also Jackson and Pettit (1990) and Walter (2010), who provide similar examples, but are not concerned with the proportionality principle.

(1) and (2) can be explained by features that are independent of the nature of causation, such as conversational pragmatics (cf. Bontly, 2005) or explanatory virtue (cf. Weslake, 2013; McDonnell, 2017; Woodward, 2021). The challenge for opponents is to argue that the data can be fully explained by these independent features.⁵ The challenge for the proponents is to argue that the proportionality requirement on causation suggested by (1)–(4) is tenable in the light of further evidence. Our focus will be on the latter challenge.

Before discussing the evidence proffered against *PROPORTIONALITY*, we should emphasize that it only provides a necessary condition on causation. In order to get at necessary and sufficient conditions, the principle is typically implemented in an existing account of causation, such as a Lewis-styled counterfactualism (e.g. List and Menzies, 2009) or interventionism (e.g. Zhong, 2020a,b). We will discuss such implementations in Section 4.2. For now we can focus on cases where the candidate causes appear to violate only the proportionality condition, if they violate any condition at all. The question we are concerned with is whether violating *PROPORTIONALITY* is enough to be eliminated as a cause. Many have argued that this is not the case.

3 The case against

Opponents of the proportionality principle argue that it is too demanding. This problem establishes itself in two ways. First, the principle appears to eliminate acceptable causes in favour of unacceptably gerrymandered phenomena. Second, intimately related phenomena sometimes share causal status, rather than excluding one another. I argue that there is a promising response to worries of the first kind, and that worries of the second kind only seriously threaten half

⁵See Maslen (2017) for some challenges for the conversational pragmatics strategy. See Schaffer (2012) for some general worries about the conversationalist pragmatics strategy when explaining the acceptability of causal claims and Gallow (2015) for some general worries about both strategies.

of PROPORTIONALITY. Let us start with the gerrymandering worry.

3.1 Gerrymandering

Many have noted that the proportionality principle seems to favour gerrymandered causes over familiar causes (e.g. McGrath, 1998; Franklin-Hall, 2016; McDonnell, 2017). To see why this is the case, consider the following example. Suppose Dougie wears a hat on weekdays, but not in the weekend. Friday morning he checks his calendar to see what day of the week it is and subsequently puts on a hat. It seems fair to say that his calendar saying it was a Friday caused him to wear a hat when leaving the apartment. However, consider the strictly stronger phenomenon [AND], which consists of both his calendar saying it was Friday *and* his having a hat available in the apartment. Had this conjunctive phenomenon not occurred, Dougie would not have left the apartment wearing a hat. What is more, if his calendar had said it was a Friday, but there was no hat available in the apartment, he would not have worn a hat when leaving. That is to say, our favoured cause does *not* screen off the target effect from [AND] and thereby fails (i). In principle, we can keep on adding required background conditions into a preposterously conjunctive phenomenon that is strictly stronger than our candidate cause and is not screened off by that cause. If unchecked, (i) favours reverse-engineered conjunctive properties over honest hardworking causes like calendar dates.

A similar problem occurs when (ii) is applied without restriction. Take the disjunctive phenomenon [OR], which consists of the calendar saying it is a Friday *or* a burglar forcing Dougie to wear a hat at gunpoint. [OR] is strictly weaker than the calendar saying it is a Friday and yet screens off his wearing a hat from the calendar saying it is a Friday. If [OR] occurs, but the calendar does not say it is a Friday, Dougie will leave his apartment wearing a hat. Generally speaking, it is easy to cook up strictly weaker phenomena that screen off respectable causes from their effects.

With the threat of both strictly stronger properties and strictly weaker properties always looming, the only phenomena that meet the proportionality requirement relative to familiar effects turn out to be horrendously unrecognizable. Only a disjunction of all the sufficient phenomena for the target effect will do (cf. Franklin-Hall, 2016, p. 572). The upshot is that PROPORTIONALITY excludes familiar causes in favour of extremely gerrymandered causes.

The proponent of PROPORTIONALITY can respond by restricting the principle to more-or-less natural phenomena (cf. Yablo, 1992a, 2003). The obvious problem with [AND] and [OR] is that they are not respectable phenomena within any credible ontology. By comparison, calendar dates and being red, while certainly not serious candidates for featuring in our fundamental theory of the universe, seem like scientifically respectable and unified phenomena. There are several ways to spell out differences in naturalness more precisely and I will not argue for one particular approach here. The important point is that some of the dominant accounts of naturalness will indeed predict that being a calendar date is more natural than [AND] and [OR]. To see this, we can look at some candidate accounts of naturalness.

One popular approach is to measure naturalness in terms of length of definition in a language that contains only names for perfectly natural properties and logical connectives (cf. Lewis, 1986b, p. 61).⁶ The perfectly natural properties will be provided by our fundamental theory of the universe, with candidates such as leptons and quarks having location and spin, and the logical connectives include disjunction and conjunction. In this language, a definition of ‘the calendar saying it is a Friday’ will be monstrously long. However, it is likely to be shorter than the definition of [AND], which will require that definition, a conjunction, *and* the definition of ‘a hat being available in the apartment’ (cf. Lewis, 1984, p. 228). *Mutatis mutandis* the same observation holds for [OR]. Consequently, ‘the calendar saying it is Friday’ comes out as more natural than

⁶See Sider (2011, Ch. 7) and Dorr (2019, Sect. 1) for discussion.

either of these gerrymandered competitors. Alternative measures of naturalness focus on the role of phenomena in scientific practice, such as biological or psychological explanation (e.g. Schaffer, 2004), or their featuring in ‘real’ patterns with other phenomena (e.g. Dennett, 1991). Here as well, ‘the calendar saying it is a Friday’ plausibly outperforms [AND] and [OR]: it patterns with it being a Friday, it explains Dougie’s elation and his trip to the liquor store, etc. By contrast, reverse-engineered phenomena tend to explain, and pattern with, primarily those effects they were reverse-engineered from, and not much else.⁷ With such approaches in hand, the defender of PROPORTIONALITY can now follow Yablo (2003, p. 324) and demand that the pursuit of a proportional cause should never lead us from more natural phenomena to less natural phenomena. Consequently, the double threat from [AND] and [OR] disappears.

When devising such a reply, the proponent of PROPORTIONALITY finds herself in good company. The distinction between gerrymandered and more-or-less natural phenomena is widely accepted as a useful philosophical tool. Well before PROPORTIONALITY entered the stage, Lewis already noted that overly disjunctive or conjunctive phenomena pose a threat for difference-making accounts of causation and proposed to eliminate them as candidate causes (1986a, p. 242). Furthermore, the ability to exclude such phenomena from our ontology in favour of more ‘natural’ or ‘joint-carving’ candidates plays an important role in theories of reference, work on natural laws, and many other areas of philosophy (cf. Dorr, 2019, Sect. 4). Advances in these areas provide the defender of PROPORTIONALITY with independent grounds for eliminating reverse-engineered phenomena like [AND] or [OR] in favour of more natural phenomena such as calendar dates.⁸ A further consequence is that gerrymandered phenomena are unlikely to settle the case against proportionality. If there is an acceptably systematic way

⁷Yablo appears to have something along these lines in mind as well when discounting gerrymandered phenomena. See for example Yablo (1992a, p. 432).

⁸Franklin-Hall (2016, p. 570) remarks that such a restriction on causal relata might be in tension with interventionist approaches to causation, as they are taken to be metaphysically non-committal. See Woodward (2018, 2021) and Blanchard (2020) for a response.

of eliminating them in favour of more natural phenomena, then they do not pose a serious problem for proponents of PROPORTIONALITY. If such an approach is not available, gerrymandered phenomena are likely to cause problems for defenders and detractors of PROPORTIONALITY alike.

3.2 Causal sharing

Unfortunately, opponents of proportionality need not rely on gerrymandered phenomena to formulate counterexamples. Sometimes non-gerrymandered, intimately related causes tolerate one another enough to share causal status. For example, Dougie's wearing a hat seems to be caused by the calendar saying it's a Friday, but also by its saying that it's a weekday — after all, he wears one *every* weekday. And Socrates dying because he drank hemlock does not appear to conflict with his dying because he drank poison (cf. Bontly, 2005; McDonnell, 2017). That is to say, all four following claims seem fine:

- (5) The calendar saying it's a weekday caused Dougie to wear his hat.
- (6) The calendar saying it's Friday caused Dougie to wear his hat.
- (7) Drinking poison caused Socrates' death.
- (8) Drinking hemlock caused Socrates' death.

In both scenarios, the strictly stronger cause is screened off from the target effect by a strictly weaker phenomenon: if it had been another weekday than Friday, Dougie would wear a hat and if it had been another poison than hemlock, Socrates would still have died. According to (ii), this screening off pattern disqualifies Friday and drinking hemlock from causing the target effect, but our judgments contradict this outcome. At least sometimes, causes are happy to share causal status with intimately related phenomena. Even if we exclude gerrymandered phenomena in favour of more natural phenomena, PROPORTIONALITY seems overly demanding.

Cases like (5)–(8) put pressure on (ii). It is harder to formulate similar counterexamples against (i). To see this consider a variation on the example of Dougie’s hat. Assume Dougie always wears his boots on Fridays, but always wears regular shoes on any other day. The following seems false:

(9) # The calendar saying it’s a weekday caused Dougie to wear boots.

Moreover, (9) seems false for exactly the reasons captured by (i): if it had been any other weekday than Friday, he would not wear his boots. (6) suggests that ‘being a weekday’ can share causal status with ‘being a Friday’, even when the latter violates (ii). If we focus on Dougie’s boots, we can see that ‘being a Friday’ does not share causal status with ‘being a weekday’ if the latter violates (i). More generally, the case of Dougie’s hat illustrates that intimately related phenomena tolerate one another’s causal status and the case of his boots illustrates that there is a limit to their tolerance. In particular, causes impose a demand on their strictly weaker companions: meet (i), or you’re out. Taken together, these cases suggest that we should stick to (i), but let go of (ii).

Of course, its being harder to formulate causal sharing counterexamples to (i), does not mean it cannot be done. As noted by Yablo (1992b, Sect. 11 & 12), there is a seemingly endless list of intimately related phenomena per cause. Some of these phenomena might appear to deliver counterexamples to (i) as well. Consider for example:

(10) Suzy throwing a rock caused the window to shatter.

(11) Suzy’s throw caused the window to shatter.

The truth of (10) seems compatible with the truth of (11), but it is at least unclear whether the window would have shattered if Suzy had thrown something without throwing a rock. Suzy throwing a rock appears to tolerate the causal status of the strictly weaker phenomena of her throw, even if the latter violates (i). So perhaps (i) is too restrictive as well.

However, it is natural to read ‘Suzy’s throw’ as referring to the phenomenon of her throwing a rock. Moreover, it is exactly this reading that explains our judgment that (11) is true. For consider how our judgments change if we check ourselves and force a less demanding reading on ‘Suzy’s throw’. Suppose we use it to pick out her making a throwing motion, regardless of whether it launched a rock, a sponge, or a marshmallow. Did *that* phenomenon cause the window to shatter? No, a stronger phenomenon — her throwing a rock — did.⁹ In short, we can allow (11) to be true by allowing ‘Suzy’s throw’ to pick out her throwing a rock. This is not a particularly costly strategy. It is well-known that we often use imprecise expressions to single out phenomena. I can refer to the intense pain in my left foot using expressions like ‘the pain in my foot’, even though this label leaves it open which foot is hurting and whether or not the pain is intense (cf. Lewis, 1986a, p. 251).¹⁰ This simple fact drastically diminishes the import of apparent counterexamples like (11). It appears that, even if all true causal claims conform with (i), they do not have to wear that conformity on their sleeve.¹¹

This same strategy is less effective against the counterexamples that plague (ii). In principle, defenders of PROPORTIONALITY could insist that ‘Socrates drinking hemlock’ picks out his drinking a lethal poison, regardless of its being

⁹Why stop there? The strictly stronger phenomenon of her throwing a rock with sufficient force in the right direction probably performs even better. Yes, but it also verges on being gerrymandered in function of the target effect, a practice we already denounced in Section 3.1. In the opposite direction, more phenomena can seem acceptable. Her throwing a heavy object and her launching a heavy object could do as causes. But note that these do not threaten her throwing a rock from being a cause according to (i), as they are strictly weaker — unless we take into account light rocks, but then we are no longer comparing intimately related phenomena. See Yablo (1992b, Sect. 11 & 12) for more on the variety of intimately related candidate causes and what can drive the selection procedure.

¹⁰See also Bennett (2002, p. 47). Maslen (2017, Sect. 5) mentions this as a possible way of addressing proportionality intuitions, but sets it aside.

¹¹We often refer to specific phenomena in even more opaque ways by quantifying over them. It is true that something Suzy did caused the window to break, because there is a phenomenon that both qualifies as her doing something and a cause of the window breaking: her throwing a rock. While existentially quantifying over phenomena will often not provide ‘just the right amount of information’, it is not in tension with the (halfway) proportionality principle as formulated here. I am indebted to Neil McDonnell and Caroline Torpe Touborg for helpful discussions on this point.

hemlock. On this reading, (8) can be true without posing problems for (ii). The problem is that claims like (8) remain true even if we force a more demanding reading on 'drinking hemlock'. Consider:

(12) Socrates drank hemlock. He didn't drink cyanide, bromide, or any other poison. Drinking *hemlock* caused him to die.

Due to the two opening sentences and the emphasis on 'hemlock', it seems clear that the 'drinking hemlock' in (12) does not pick out the less specific event of drinking some lethal poison or other. It quite explicitly picks out drinking hemlock, and not any other poison. Even so, (12) seems to contain an acceptable causal claim: Drinking hemlock caused Socrates to die. The 'blaming the labels' strategy is now unavailable, and the defender of full-blown proportionality still has a counterexample on her hands.

For contrast, consider a similar attempt to inoculate the 'blaming the labels' strategy when used to defend (i):

(13) # It could have been a Monday, Tuesday or Wednesday instead of a Friday. It's being a *weekday* caused Dougie to wear his boots.

The preceding sentence and the focus on 'weekday' emphasize that 'it's being a weekday' is taken to pick out exactly its being a weekday and not a more specific phenomenon. As opposed to (12), (13) does not appear to contain an acceptable causal claim. If Dougie only wears boots on Fridays, then it seems false to say that it's being a weekday caused him to wear his boots. Rather than inoculating the 'blaming the labels' strategy, examples in which one forces the weaker reading appear to corroborate (i).

If we now sum up the case against PROPORTIONALITY, it appears that only one of the two subprinciples is under serious threat. True, gerrymandered phenomena pose a threat to both (i) and (ii), but there are good reasons to believe that such phenomena can be eliminated from the causal competition on independent grounds. When we look at counterexamples containing more-or-less

natural phenomena, they appear to primarily threaten (ii). Moreover, apparent counterexamples to (i) can be explained away by relying on the familiar fact that we sometimes pick out more specific phenomena than is suggested by the labels we use. In light of this evidence, it appears that defenders of proportionality could meet their critics halfway, and endorse a weakened proportionality principle that retains (i), but drops (ii).

4 Halfway Proportionality

A closer look at the case against PROPORTIONALITY suggested that one half of the principle withstands the standard criticisms. We can meet the opponents halfway by adopting the following principle:

Halfway Proportionality P causes E only if

- (i) For any $P+$, strictly stronger than P , P screens off E from $P+$

This halfway principle can be motivated further with recent work on the significance of causal cognition.

Hitchcock (2012) notes that empirical and philosophical work on causation has converged on the thesis that we are interested in causes because the correlation with their effects is typically *insensitive* to further factors in a way that non-causal correlations are not. For example, the positive correlation between me drinking coffee and me feeling chipper ten minutes later is (largely) insensitive to changes in time of day and day of the year. This insensitivity to further factors makes causal correlations reliably exportable to a broader variety of scenarios than non-causal correlations. One can reliably predict or affect when I'll be chipper across a wide variety of scenarios by predicting or affecting when I have my cup of coffee in those scenarios. By contrast, the correlation between me drinking coffee and the sun rising ten minutes later is sensitive to changes in

when and where I drink that cup. Consequently, one can reliably predict the sun coming up in ten minutes based on my drinking coffee only in those scenarios where I happen to drink a cup of coffee ten minutes before sunrise.

Such (in)sensitivity to further factors illuminates several issues in philosophy of causation. Woodward (2006) argues that it draws a distinction between good and bad cases of absence causation. My not taking care of my plants is plausibly a cause of their death, but Jeff Bezos not taking care of them is not seen as a cause. Correspondingly, the correlation between my lack of care and their death is relatively insensitive to normal changes in the background conditions, whereas Jeff Bezos' lack of care towards my plants only result in their decay in those scenarios where I fail to take care of them.¹² Lombrozo (2010) argues that insensitivity considerations explain the importance of mechanisms connecting cause and effect, as correlations between phenomena that are connected by a mechanism tend to be insensitive to changes in their surroundings. Finally, even the crucial distinction between causation and correlation confounded by a common cause can be understood in terms of insensitivity. The correlation between yellowed teeth and lung cancer is sensitive to whether the yellowness is caused by smoking or coffee, but the correlation between smoking and lung cancer is largely insensitive to how the smoking is caused (cf. Woodward, 2014, p. 710). According to all of these accounts, insensitivity and causation go hand in hand.

Now note that the correlation between phenomena that violate only (ii) and their target effect is insensitive in a way that correlations that violate only (i) is not. In particular, correlations that violate (i) are sensitive to changes in the realization of the first phenomenon, whereas correlations that fail (ii) do not suffer from the same defect. For example, being coloured is followed by Sophie pecking only in those scenarios where the being coloured is realized by a shade of

¹²Admittedly, my not taking care of my plants would not result in their decay if Jeff Bezos decided to take care of them — he is a man of unlimited resources. However, such a change in the background conditions hardly qualifies as 'normal'.

red. By contrast, being scarlet does result in Sophie pecking across a wide variety of scenarios, as she is trained to peck at all red objects.¹³ Consequently, the scarlet-pecking correlation is reliably exportable to a wide variety of scenarios, whereas the coloured-pecking correlation is not. According to recent work on causation and (in)sensitivity, this suggests that the scarlet-pecking correlation is causal, but the coloured-pecking correlation is not.

HALFWAY PROPORTIONALITY fits with cases like Dougie's, has a promising strategy of dealing with counterexamples, and is supported by recent work on causation. Consequently, I submit that it has a better chance at survival than PROPORTIONALITY. In the remainder of this section, I discuss the consequences of adopting HALFWAY PROPORTIONALITY rather than adopting PROPORTIONALITY or no proportionality-like principle at all. First, I focus on the impact on the exclusion debate. Then I will discuss the adjustments HALFWAY PROPORTIONALITY imposes on standard difference-making theories of causation.

4.1 Halfway Proportionality and Causal Exclusion

PROPORTIONALITY'S central selling point is that it provides a straightforward solution to so-called exclusion arguments.¹⁴ It would certainly be good news if HALFWAY PROPORTIONALITY can retain this feature. In this section, I argue that, while our halfway principle does not deliver the same result, it provides a significant contribution to the debate on causal exclusion nonetheless.

The problem of causal exclusion threatens the causal status of *higher-level* phenomena. Higher-level phenomena are those phenomena that are not indi-

¹³This is not to say that phenomena that meet (i) relative to an effect will be insensitive to *all* changes in non-gerrymandered strictly stronger phenomena. For example, an object being red or scarlet will have some possible realizations with thermodynamically abnormal futures. Perhaps the object ejects a particle with sufficient acceleration to kill Sophie on the spot, seriously injuring her chances of pecking at it. As these complications are well-discussed in the literature on counterfactuals (e.g. Sandgren and Steele, 2020; Stefánsson, 2018) and causation (e.g. Loewer, 2007a; Albert, 2015; Woodward, 2018), I set them aside here.

¹⁴See Yablo (1992b,a, 1997); List and Menzies (2009); Shoemaker (2001); Raatikainen (2010) and Zhong (2020a,b).

viduated in terms of our most fundamental theory, but in more coarse-grained terms, such as speeding past speedometers, hurricanes occurring or objects being red. Using our terminology, we can say that higher-level phenomena are strictly weaker than the fundamental physical phenomena underlying them, which are individuated in terms of leptons and quarks having a certain mass and spin.¹⁵

Exclusion problems arise because it seems more and more likely that all effects have physically sufficient physical conditions at any time preceding them (Papineau, 2001). If these physical conditions qualify as causes, then each effect has a sufficient physical cause at any time preceding it. With leptons and quarks doing all the causing required for any effect you like, there appears to be no room for our familiar higher-level causes like redness, speeding past speedometers or the desire for a sip of water. We can represent the two principles driving this reasoning as follows (cf. Yablo, 1992a, p. 429):

Causal Closure Every effect has a sufficient physical cause at any given time t .

Causal Exclusion For any two non-identical co-occurrent phenomena A and B , and any effect E : if A is a sufficient cause for E , then B does not cause E .¹⁶

Taken together, these principles entail that effects cannot have causes that are non-identical to physical phenomena. The upshot is that our everyday and scientific causal talk, which mostly concerns higher-level causes, is systematically mistaken.

PROPORTIONALITY provided a straightforward solution to this problem. For if

¹⁵I'm assuming physicalism here, but this is merely for convenience. The reader can substitute fundamental physics with her favoured fundamental theory.

¹⁶Exclusion principles usually contain a disclaimer like 'unless it is a case of genuine overdetermination', which is balanced by a separate 'there is no systematic genuine overdetermination' principle (e.g. Bennett, 2007). To keep things simple, we will set issues about genuine overdetermination aside here.

PROPORTIONALITY is true, then CAUSAL CLOSURE is probably false. Even if we admit that, for every familiar effect, there is a physical condition that is physically sufficient for its occurrence, it is unlikely that such conditions meet (ii) relative to those effects, as they are rife with irrelevant detail. For example, consider the physically sufficient condition for Sophie's pecking the scarlet pebble a second before her doing so. Call that phenomenon $P+$. According to contemporary physics, $P+$ will have to span a sphere with a radius of 300 000 km surrounding the pecking's location and will have to be maximally fine-grained, in that it fixes the exact mass, spin and position of all the fundamental particles in that sphere (cf. Loewer, 2007b). There is a multitude of co-occurrent phenomena that are strictly weaker than $P+$ and screen off the pecking from the occurrence of $P+$. For example, consider a phenomenon that is identical to $P+$, except that it only requires the pebble to be red, rather than specifying its physical constitution in minute detail. This phenomenon will screen off the pecking from $P+$. That is to say, if this phenomenon had occurred without $P+$ occurring, Sophie would still have pecked. On the plausible assumption that these observations generalize, physically sufficient conditions do not qualify as causes of familiar effects because they systematically violate (ii) relative to those effects.

So (ii) could save the day. With CAUSAL CLOSURE out of the way, CAUSAL EXCLUSION stops threatening all higher-level causes. Unfortunately, (ii) is exactly the clause we let go off earlier. Moreover, (i), the clause we decided to stick to in HALFWAY PROPORTIONALITY, does nothing to contradict CAUSAL CLOSURE. After all, the threat to phenomena like $P+$ came from strictly weaker phenomena, and (i) does not entail that strictly weaker phenomena can exclude their strictly stronger companions.¹⁷ One might thus worry that we have are holding on to the useless half of the proportionality principle.

The situation is not quite as bad. True, HALFWAY PROPORTIONALITY does not

¹⁷Of course, one could reject $P+$ as a cause on the grounds of its being too gerrymandered, but then proportionality considerations are not playing a decisive role.

resolve exclusion problems on its own, but the case we made in its favour hints at a solution nonetheless. In the absence of (ii), (i) allows for intimately related phenomena to cause the same effect. That is exactly what is going on in the examples which drove us to adopt HALFWAY PROPORTIONALITY instead of old-fashioned PROPORTIONALITY: Dougie's wearing a hat being caused by the calendar saying it is a Friday, does not conflict with its being caused by the calendar saying it is a weekday. Consequently, our case in favour of HALFWAY PROPORTIONALITY suggests that CAUSAL EXCLUSION is false. Even if all effects have physically sufficient physical causes, this would be compatible with there being co-occurrent, non-identical, but intimately related causes of those effects. For example, *P+* causing Sophie to peck would be compatible with the strictly weaker phenomenon of the pebble being red also causing Sophie to peck. Of course, *suggesting* a solution is not the same as fully motivating it, and a proper defense of such a 'compatibilism' about intimately related causes would require more arguments than I can provide here. However, forceful defenses of this compatibilism can readily be found in Bennett (2003, 2008); Schaffer (2003); Sider (2003); Stoljar (2008) and Woodward (2008, 2015), among others. Moreover, there are independent reasons to be wary of addressing the exclusion problem by focusing on the causal status of physically sufficient conditions (cf. Vaassen, 2021a). The solution suggested by HALFWAY PROPORTIONALITY and the examples that motivate it might require more legwork than the one delivered by the full-blown proportionality principle, but there is reason to think it is the right one. And indeed, some recent accounts of higher-level causation conform with (i) without committing to (ii) (cf. Woodward, 2018, 2020).

Moreover, HALFWAY PROPORTIONALITY retains a feature of PROPORTIONALITY that is both interesting and controversial: higher-level causes have to earn their keep. Opponents of proportionality sometimes maintain that phenomena automatically inherit the causal status of their underlying strictly stronger phenomena (e.g. Shapiro and Sober, 2007; Shapiro, 2010). If this is true, then

CAUSAL EXCLUSION is false. However, cases like Dougie's boots have been used to demonstrate that there cannot be such universal causal inheritance (e.g. Jackson and Pettit, 1990), and some have taken this to mean that exclusion problems remain (e.g. Walter, 2010). HALFWAY PROPORTIONALITY carves out a promising middle road: phenomena *can* inherit causal status from their underlying strictly stronger phenomena, but only if they meet (i). That is to say, CAUSAL EXCLUSION is false, but merely standing in an intimate relation to the sufficient cause of (or a physically sufficient condition of) an effect does not guarantee causal status relative to that effect either. Instead, higher-level causes have to earn their keep by screening off their effects from the strictly stronger phenomena underlying them.

This struggle for causal status means that we can, at least in principle, empirically test whether higher-level phenomena are outcompeted by their strictly stronger companions. On the assumption that my desire for a sip of water is metaphysically necessitated by a strictly stronger physical phenomenon that causes my reaching for a glass, the question remains whether I would still have reached for the glass if the desire occurred without this strictly stronger physical phenomenon. Many have argued that this is almost certainly the case (e.g. Campbell, 2008, 2010), and a brief moment of reflection supports this judgment. Surely, I would have reached for the glass even if some leptons and quarks had differed ever so slightly in position or spin. This does not yet establish the causal status of my desire though. There are plenty of phenomena that are strictly stronger than my desire, but strictly weaker than the maximally specific stronger physical phenomenon. Such intermediate phenomena are likely to include chemical, biological and neurological phenomena. According to (i), my desire should screen off my reaching from all these intermediate phenomena. At least in principle, we can test whether the correlation between my desire for water and my reaching for it is insensitive to changes in any of these stronger phenomena. Contrary to what opponents of proportionality have claimed (e.g.

Baumgartner, 2018), finding out whether or not higher-level phenomena like desires are causes might be settled by empirical research.¹⁸

To sum up, *HALFWAY PROPORTIONALITY* does not solve the causal exclusion problem, but it makes three distinct contributions to the exclusion debate. First, the arguments in its favour suggest that *CAUSAL EXCLUSION* is false. Second, it predicts that causal inheritance is possible, but not universal. Third, it predicts that the causal status of higher-level phenomena is empirically testable. Moreover, the extra work required to motivate the rejection of *CAUSAL EXCLUSION* has already been done by compatibilists. While falling short of a cure-all for exclusion worries, *HALFWAY PROPORTIONALITY* provides a nuanced and interesting picture of higher-level causation.

4.2 Halfway proportionality and difference-making

HALFWAY PROPORTIONALITY also has repercussions for some standard theories of causation. Just like *PROPORTIONALITY*, the principle entails that simple counterfactual dependence or ‘difference-making’ does not suffice for causation.¹⁹ In this Section, I discuss how some popular difference-making accounts of causation fail the principle and what kind of changes *HALFWAY PROPORTIONALITY* imposes.

¹⁸Of course, some philosophical questions will still need resolving as well. For example, it matters whether the necessitation relation between higher-level or mental phenomena and their underlying stronger phenomena is of *metaphysical* strength or *nomological* strength. There are strong reasons to believe that, if *CAUSAL CLOSURE* is true, phenomena that are merely nomologically necessitated cannot be causes (cf. Bennett, 2008; Vaassen, 2021b). Though see Kroedel (2015, 2020), List and Stoljar (2017), and Vaassen (2019) for strategies to resist this conclusion.

¹⁹I restrict myself to difference-making accounts of causation, as these are more commonly associated with proportionality. Note however, that process theories of causation such as Dowe (2000) appear to conflict with *HALFWAY PROPORTIONALITY* as well. On the assumption that my sweater’s being red and its being scarlet are co-located and realized by the same physical phenomenon, one can transfer a conserved quantity on a target effect if and only if the other one does. So on process theories their causal status stands or falls together. This clearly contradicts *HALFWAY PROPORTIONALITY*. Consequently, adding a process requirement to difference-making accounts, as proposed by Schaffer (2001) and Loew (2019), will not save such accounts from counterexamples like Dougie’s boots either.

The idea behind difference-making accounts of causation is that effects counterfactually depends on the cause. If the cause had not occurred, the effect had not occurred either. This idea has gained traction via the account proposed by Lewis (1973a, 1979). We can represent the sufficiency conditions comprised in Lewis' proposal with the following two counterfactual tests:

(C) For any two phenomena, C and E , the occurrence of C causes the occurrence of E if

(+) C occurs $\Box \rightarrow E$ occurs, and

(-) C does not occur $\Box \rightarrow E$ does not occur.

where the semantics of $\Box \rightarrow$ are such that

For any two propositions P and Q , $P \Box \rightarrow Q$ is true *iff* there is a possible world where P and Q hold which is closer to the actual world than any possible world where P and not- Q hold (or there are no possible worlds where P holds).

Our cases violating HALFWAY PROPORTIONALITY clearly meet the Lewisian criteria for causing the target effect. If we present Sophie with a red pebble and she pecks, the pebble's being coloured passes both counterfactual tests relative to Sophie's pecking. The closest possible world where the pebble is coloured is the actual one, and Sophie pecks in the actual world, as per (+). As per (-), Sophie will not peck in nearby possible worlds where the pebble is not coloured, because non-coloured pebbles are not red, and Sophie pecks exclusively at red things. If HALFWAY PROPORTIONALITY is true, the standard Lewisian account fails to provide sufficient conditions for causation.

List and Menzies (2009) demonstrated that the Lewisian account can be made to capture condition (i) by adopting a *weak centering* thesis. According to the weak centering thesis, there are non-actual worlds that are equally close to actuality as the actual world. For example, in the Sophie case, we can consider

the world in which the pebble is coloured, but not red, as equally close to actuality as the actual world. Consequently, we need to consider this world as well to evaluate (+). If we do this, it comes out as false, because in nearby worlds where the pebble is coloured non-red, Sophie does not peck. By contrast, (+) remains true if we look at the closest worlds where the pebble is red, but differs in its shade of redness from actuality. Consequently, redness comes out as a cause of Sophie's pecking and being coloured does not, just as HALFWAY PROPORTIONALITY required.²⁰

Weak centering forces us to look at worlds where the purported cause differs slightly in its realization when evaluating (+). Consequently, this counterfactual tracks the insensitivity to realization that we wanted to capture with HALFWAY PROPORTIONALITY. Note however, that adopting weak centering is not required to achieve this result. In principle, one could insist that these nearby world with different realizers are relevant to causal claims, without insisting that these worlds are as close to actuality as the actual world. One could do so by simply adding condition (i) to the Lewisian account instead of adopting the weak centering thesis. In the end, what we require is an insensitivity-to-realization-tracking criterion in our account of causation.

Crucially, such an insensitivity-tracking criterion is *not* present in some of the popular contemporary difference-making accounts of causation. To see this, we can look at two examples: contrastivism (e.g. Schaffer, 2005) and interventionism (e.g. Woodward, 2003, 2021).

²⁰An anonymous referee pointed out that one might worry that adopting List and Menzies' framework will eventually commit one to full-blown PROPORTIONALITY. Note however that merely adding weak centering to (C) does not entail (ii). Even if we consider nearby worlds that slightly differ from actuality, being scarlet still passes (+) as all possible worlds where the pebble is scarlet are trivially worlds where it is red as well. Establishing (ii) requires the extra step of maintaining that the closest relevant non-scarlet world is still a world where the pebble is red, such that 'being scarlet' fails (–) (cf. List and Menzies, 2009, p. 488). As noted by Woodward (2015, fn. 1) and Zhong (2020b, p. 306) this is a controversial step and it goes against Lewis' advice to "wholly excise" the phenomena one eliminates to assess counterfactual dependence (e.g. Lewis, 2000, p. 190). If one adopts weak centering, but still follows Lewis' advice on excision, one can accommodate (i) without committing to (ii) See McDonnell (2019) for an in-depth discussion on event excision and difference-making accounts of causation.

According to contrastivism, the Lewisian account needs to be complemented with *contrast classes* for both the cause and the effect. These contrast are typically hidden and provided by the conversational context. The general idea is that making the role of these hidden contrast classes more explicit helps to address some problem cases in philosophy of causation. However, considering the contrast classes does not always resolve the problems with claims violating (i). Consider again Dougie's boots. This time with contrasts made explicit:

(14) # The calendar saying it is weekday rather than the weekend caused Dougie to wear his boots rather than his regular shoes.

By the contrastivist's lights (14) should still come out as true. We should check whether Dougie wears regular shoes in nearby worlds where the calendar says it is a Saturday or a Sunday. Based on what we know about the case, this counterfactual still comes out as true. Contrast classes do not appear to capture what is wrong with the case of Dougie's boots and the contrastivist is still in need of an extra criterion such as HALFWAY PROPORTIONALITY.²¹

According to interventionism, causal relations can be characterized in terms of what happens under certain interventions. The causal relata are represented as variables that can take at least two values and the central question is whether *interventions* on the value of the cause variable correlate with changes in the value of the effect variable. Where an intervention is a causal interaction that leaves all other relevant variables unchanged. Setting aside some complications, we can represent the account as follows (Woodward, 2021, p. 77):

(TC) *C* (a variable representing the putative cause) causes *E* (a variable representing the effect) in background circumstances *B* iff there is some possible intervention that changes the value of *C*

²¹In fairness, the introduction of contrast classes was not intended to address cases like Dougie's boots. Schaffer's central papers developing contrastivism do not address the issue of proportionality (Schaffer, 2005, 2012) and contrastivist proposals to address surrounding issues, such as Weslake (2013) and Maslen (2017) focus on (ii) rather than (i).

such that if that intervention were to occur in B , there would be an associated change in the value of E or in the probability distribution $P(E)$ of those values.

Woodward (2021, Ch. 8) provides an in-depth discussion of why this definition does not embody proportionality-like requirements. For our purposes, the central point is that it does not capture HALFWAY PROPORTIONALITY. Changing the date from a Friday to a Saturday counts as an intervention on the calendar saying it is a weekday, and we know that such a change associates with Dougie's wearing rather regular shoes rather than boots.²² If HALFWAY PROPORTIONALITY holds, interventionist definitions of causation, such as (TC), require revision.

One way to revise (TC) in accordance with HALFWAY PROPORTIONALITY is to demand that the intervention-supporting association between the cause and effect variables remains robust under changes in the variables representing strictly stronger phenomena.²³ For example, the association between the calendar saying it is a weekday and Dougie's wearing a hat is robust under changes in which specific day of the week the calendar says it is, whereas the association between the calendar saying it is a weekday and his wearing boots is sensitive to these changes. We can leave the specifics of how to implement such a requirement more precisely for future research. The central point is that such an insensitivity-tracking criterion is required to capture HALFWAY PROPORTIONALITY and that such a criterion is not present in the standard interventionist definition.

²²One might worry that this change cannot leave the value of the variable representing what specific day of the week it is according to the calendar, but such intimately related variables are exempted from the holding fixed requirement (cf. Woodward, 2008, 2015).

²³Zhong (2020a, p. 302–304) makes a similar proposal to capture (i), but focuses on binary variables.

5 Conclusion

The proportionality principle divides philosophers writing on causation. Some take it to be a silver bullet against exclusion problems. Others have pointed out that the principle meets plenty of counterexamples. Unfortunately, the clause that is typically targeted by the critics is also the clause that did all the heavy lifting in resolving the exclusion problem. Even so, the less discussed clause is significant enough in its own right. It provides a natural fit with compatibilist responses to exclusion arguments and imposes further adjustments on difference-making accounts of causation. This clause also has more resources available to explain away apparent counterexamples. Consequently, I propose that we reject full-blown PROPORTIONALITY and adhere to HALFWAY PROPORTIONALITY instead.

Acknowledgements

This material was originally presented at the *Causal Distinctions* workshop at the University of Cologne and the Higher Seminar at Umeå University. I benefited greatly from the questions and comments from the participants there. Special thanks are due to Martín Abreu Zavaleta, Ethan Nowak, Alex Sandgren, and two anonymous referees for excellent comments on the written text, as well as to Jim Woodward for sending preprints of his book and for swiftly answering my questions on (halfway) proportionality.

References

- Albert, D. Z. (2015). *After Physics*. Harvard University Press.
- Baumgartner, M. (2018). The inherent empirical underdetermination of mental causation. *Australasian Journal of Philosophy*, 96(2):335–350.

- Bennett, J. (2002). What events are. In Gale, R. M., editor, *The Blackwell Guide to Metaphysics*, pages 43–65. Blackwell.
- Bennett, K. (2003). Why the exclusion problem seems intractable, and how, just maybe, to tract it. *Noûs*, 37(3):471–497.
- Bennett, K. (2007). Mental causation. *Philosophy Compass*, 2(2):316–337.
- Bennett, K. (2008). Exclusion again. In Kallestrup, J. and Hohwy, J., editors, *Being Reduced*, pages 280–307. Oxford University Press.
- Blanchard, T. (2020). Explanatory abstraction and the goldilocks problem: Interventionism gets things just right. *British Journal for the Philosophy of Science*, 71(2):633–663.
- Bontly, T. D. (2005). Proportionality, causation, and exclusion. *Philosophia*, 32(1):331–348.
- Cameron, R. P. (2010). How to have a radically minimal ontology. *Philosophical Studies*, 151(2):249–264.
- Campbell, J. (2008). Interventionism, control variables and causation in the qualitative world. *Philosophical Issues*, 18(1):426–445.
- Campbell, J. (2010). Control variables and mental causation. *Proceedings of the Aristotelean Society*, 110(October 2009):15–30.
- Dennett, D. C. (1991). Real patterns. *Journal of Philosophy*, 88(1):27–51.
- Dorr, C. (2019). Natural Properties. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2019 edition.
- Dowe, P. (2000). *Physical Causation*. Cambridge University Press.

- Franklin-Hall, L. R. (2016). High-level explanation and the interventionist's 'variables problem'. *British Journal for the Philosophy of Science*, 67(2):553–577.
- Gallow, J. D. (2015). The emergence of causation. *Journal of Philosophy*, 112(6):281–308.
- Hitchcock, C. (2012). Portable causal dependence: A tale of consilience. *Philosophy of Science*, 79(5):942–951.
- Hoffmann-Kolss, V. (2014). Interventionism and Higher-level Causation. *International Studies in the Philosophy of Science*, 28(1):49–64.
- Jackson, F. and Pettit, P. (1990). Causation and the philosophy of mind. *Philosophy and Phenomenological Research*, 50(n/a):195–214.
- Kroedel, T. (2015). Dualist mental causation and the exclusion problem. *Noûs*, 49(2):357–375.
- Kroedel, T. (2020). *Mental Causation: A Counterfactual Theory*. Cambridge University Press.
- Lewis, D. (1973a). Causation. *The Journal of Philosophy*, 70(17):556–567.
- Lewis, D. (1973b). *Counterfactuals*. Oxford University Press.
- Lewis, D. (1979). Counterfactual dependence and time's arrow. *Noûs*, 13(4):455–476.
- Lewis, D. (1984). Putnam's paradox. *Australasian Journal of Philosophy*, 62(3):221–236.
- Lewis, D. (1986a). Events. In Lewis, D., editor, *Philosophical Papers Vol. II*, pages 241–269. Oxford University Press.
- Lewis, D. (1986b). *On the Plurality of Worlds*. Wiley-Blackwell.

- Lewis, D. (2000). Causation as influence. *Journal of Philosophy*, 97(4):182–197.
- List, C. and Menzies, P. (2009). Non-reductive physicalism and the limits of the exclusion principle. *The Journal of Philosophy*, 106(9):475–502.
- List, C. and Stoljar, D. (2017). Does the exclusion argument put any pressure on dualism? *Australasian Journal of Philosophy*, 95(1):96–108.
- Loew, C. (2019). Causes as differencemakers for processes. *Philosophy and Phenomenological Research*, 98(1):89–106.
- Loewer, B. M. (2007a). Counterfactuals and the second law. In Price, H. and Corry, R., editors, *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, pages 293–326. Oxford University Press.
- Loewer, B. M. (2007b). Mental causation, or something near enough. In McLaughlin, B. P. and Cohen, J. D., editors, *Contemporary Debates in Philosophy of Mind*, pages 243–64. Blackwell.
- Lombrozo, T. (2010). Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology*, 61(4):303–332.
- Maslen, C. (2017). Pragmatic explanations of the proportionality constraint on causation. *Making a Difference: Essays on the Philosophy of Causation*. Oxford University Press, New York, pages 58–72.
- McDonnell, N. (2017). Causal exclusion and the limits of proportionality. *Philosophical Studies*, 174(6):1459–1474.
- McDonnell, N. (2019). The non-occurrence of events. *Philosophy and Phenomenological Research*, 99(2):269–285.
- McGrath, M. (1998). Proportionality and mental causation: A fit? *Philosophical Perspectives*, 12:167–176.

- Papineau, D. (2001). The rise of physicalism. In Gillett, C. and Loewer, B. M., editors, *Physicalism and its Discontents*, pages 3–36. Cambridge University Press.
- Paul, L. A. (2006). Coincidence as overlap. *Noûs*, 40(4):623–659.
- Paul, L. A. (2007). Constitutive overdetermination. In Campbell, J. K., O’Rourke, M., and Silverstein, H. S., editors, *Causation and Explanation*, pages 265–290. MIT Press.
- Pereboom, D. (2002). Robust non-reductive materialism. *Journal of Philosophy*, 99(10):499–531.
- Raatikainen, P. (2010). Causation, exclusion, and the special sciences. *Erkenntnis*, 73(3):349–363.
- Sandgren, A. and Steele, K. (2020). Levelling counterfactual scepticism. *Synthese*, pages 1–21.
- Schaffer, J. (2001). Causes as probability raisers of processes. *Journal of Philosophy*, 98(2):75–92.
- Schaffer, J. (2003). Overdetermining causes. *Philosophical Studies*, 114(1-2):23–45.
- Schaffer, J. (2004). Two conceptions of sparse properties. *Pacific Philosophical Quarterly*, 85(1):92–102.
- Schaffer, J. (2005). Contrastive causation. *Philosophical Review*, 114(3):327–358.
- Schaffer, J. (2012). Causal contextualisms. In Blaauw, M., editor, *Contrastivism in Philosophy: New Perspectives*, pages 43–71. Routledge.
- Schaffer, J. (2015). What not to multiply without necessity. *Australasian Journal of Philosophy*, 93(4):644–664.

- Shapiro, L. (2010). Lessons from causal exclusion. *Philosophy and Phenomenological Research*, 81(3):594–604.
- Shapiro, L. and Sober, E. (2007). Epiphenomenalism. The do's and don'ts. In Wolters, G. and Machamer, P., editors, *Thinking about causes: From Greek philosophy to modern physics*, pages 235–264. University of Pittsburgh Press.
- Shapiro, L. and Sober, E. (2012). Against proportionality. *Analysis*, 72(1):89–93.
- Shoemaker, S. (2001). Realization and mental causation. In Gillett, C. and Loewer, B. M., editors, *Physicalism and its Discontents*, pages 74–98. Cambridge University Press.
- Sider, T. (2003). Review: What's so bad about overdetermination? *Philosophy and Phenomenological Research*, 67(3):719–726.
- Sider, T. (2011). *Writing the Book of the World*. Oxford University Press.
- Stefánsson, H. (2018). Counterfactual skepticism and multidimensional semantics. *Erkenntnis*, 83(5):875–898.
- Stoljar, D. (2008). Distinctions in distinction. In Kallestrup, J. and Hohwy, J., editors, *Being Reduced: New Essays on Causation and Explanation in the Special Sciences*, pages 263–279. Oxford University Press.
- Vaassen, B. (2019). *Causal After All : A Model of Mental Causation for Dualists*. PhD thesis.
- Vaassen, B. (2021a). Causal exclusion without causal sufficiency. *Synthese*, 198:10341–10353.
- Vaassen, B. (2021b). Dualism and exclusion. *Erkenntnis*, 86(3):543–552.
- Walter, S. (2010). Taking realization seriously: No cure for epiphobia. *Philosophical Studies*, 151(2):207–226.

- Weslake, B. (2013). Proportionality, contrast and explanation. *Australasian Journal of Philosophy*, 91(4):785–797.
- Woodward, J. (2003). *Making Things Happen: a Theory of Causal Explanation*. Oxford University Press.
- Woodward, J. (2006). Sensitive and insensitive causation. *Philosophical Review*, 115(1):1–50.
- Woodward, J. (2008). Mental causation and neural mechanisms. In Kallestrup, J. and Hohwy, J., editors, *Being Reduced*, pages 218–263. Oxford University Press.
- Woodward, J. (2014). A functional account of causation; or, A defense of the legitimacy of causal thinking by reference to the only standard that matters—usefulness. *Philosophy of Science*, 81(5):691–713.
- Woodward, J. (2015). Interventionism and causal exclusion. *Philosophy and Phenomenological Research*, 91(2):303–347.
- Woodward, J. (2018). Explanatory autonomy: The role of proportionality, stability, and conditional irrelevance. *Synthese*, pages 1–29.
- Woodward, J. (2020). Causal complexity, conditional independence, and downward causation. *Philosophy of Science*, 87(5):857–867.
- Woodward, J. (2021). *Causation With A Human Face: Normative Theory and Descriptive Psychology*. Oxford University Press.
- Yablo, S. (1992a). Cause and essence. *Synthese*, 93(3):403–449.
- Yablo, S. (1992b). Mental causation. *Philosophical Review*, 101(2):245–280.
- Yablo, S. (1997). Wide causation. *Philosophical Perspectives*, 11:251–281.
- Yablo, S. (2003). Causal Relevance. *Philosophical Issues*, 13(1):316–328.

Zhong, L. (2014). Sophisticated exclusion and sophisticated causation. *Journal of Philosophy*, 111(7):341–360.

Zhong, L. (2020a). Intervention, fixation, and supervenient causation. *Journal of Philosophy*, 117(6):293–314.

Zhong, L. (2020b). Taking emergentism seriously. *Australasian Journal of Philosophy*, 98(1):31–46.