Asian Journal of Philosophy (2024) 3:54
https://doi.org/10.1007/s44204-024-00184-5

ARTICLE SYMPOSIUM

Against causal arguments in metaphysics

Bram Vaassen[1]

Received: 29 April 2024 / Accepted: 11 August 2024 / Published online: 2 September 2024

Abstract
Traditionally, causal arguments for physicalism have been taken to favour a 'reductive' brand of physicalism, according to which all the mental stuff is identical to some of the physical stuff. Many flaws have been found with these traditional causal arguments. Zhong (*Asian Journal of Philosophy, 2*(2), 1–9, 2023) develops a new causal argument that avoids these flaws and favours a milder, non-reductive brand of physicalism instead. The conclusion is that all mental stuff is metaphysically necessitated by some of the physical stuff. I argue that neither the traditional nor the new causal argument holds much sway over non-physicalism. The problem is that causation just does not run that deep. It is a fairly superficial relationship and a poor guide to the metaphysically weighty facts of our world, such as what is identical to what, and what metaphysically necessitates what.

Keywords Mental causation · Causal arguments · Dualism · Physicalism · Exclusion arguments · Lei Zhong

1 Introduction

I am a physicalist by temperament. Someone tells me that zombies are conceivable, and therefore, our pain is metaphysically separable from all physical goings-on, and my mind immediately goes 'I don't know about any of *that*…' The same goes for Mary in her black-and-white room, spectrum inversions, and our inability to imagine what it is like to be a bat. I am quite compelled by the premises. I feel like I can conceive of zombies and of the possibility that my seeing red is phenomenally like what it is like for you to see blue. I can confidently imagine brilliant colour scientists without colour experience, but cannot fathom what it is like to echolocate. It's just that I'm reluctant to draw the proposed conclusion. The main tenet of physicalism, that pains, desires, and beliefs are at bottom made up out of the same physical

✉ Bram Vaassen
bram.vaassen@umu.se

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

building blocks as cars, iPhones, and hurricanes, is something my mind has a hard time letting go off.

That said, I have started to feel reluctant to accept the arguments favouring physicalism as well. I'd like nothing more than to embrace the conclusion; I'm just not certain that the arguments take me all the way there. Zhong (2023) rekindles hope for temperamental physicalists like me. He argues that the much-discussed causal argument for physicalism can be reformulated to withstand much of the criticism it has received lately. And indeed, Zhong's revised causal argument is a clear improvement over its predecessors. Even so, it leaves me unconvinced that we have an airtight case against non-physicalism. The problem with the argument, as with previous iterations of it, is that it relies on causation. According to our best theories, causation is a metaphysically superficial dependence relation. While such dependence relations are of crucial importance for scientific endeavours and everyday life, they do not serve as a guide for the identity relations or metaphysical necessitation relations that the physicalist is after.

I will spell out my concerns as follows. First, I provide a brief characterization of the standard causal argument, its problems, and Zhong's reformulation (§2). Second, I spell out my general concerns with using causal arguments for physicalism (§3). Third, I show how these concerns affect Zhong's reformulated argument (§4). I conclude that we should re-evaluate the use of causal arguments in metaphysical debates.

2 Causal arguments

By far the most influential argument favouring physicalism is the so-called exclusion argument (Gebharter, 2017; Kim, 1989, 2005; Ney, 2009), also known as the causal argument (Papineau, 2001, 2002). It starts from the premise that all physical effects are necessitated by their physical history and the fundamental laws of physics, and invites us to conclude that, unless pains never cause physical effects like shrieking, all mental stuff must be identical to some of the physical stuff. Zhong formulates the argument as follows (Zhong, 2023, p. 3):

1a) [Interaction] Mental properties are causally relevant to some physical effects.
1b) [Closure] Every physical effect has a sufficient physical cause (at any time at which it has a cause).
1c) [Exclusion] There is no systematic overdetermination by physical and non-physical causes.
1d) [Identity] Therefore, mental properties are physical properties.

Many have responded that the argument, in this form, builds on an overly demanding notion of causation (Block, 2003; Kroedel, 2020; Russo, 2016). Zhong (2023) argues that this defect can be fixed if one is willing to settle for the conclusion that all mental stuff is metaphysically necessitated by the physical stuff, rather than being identical with it. We could thus get a causal argument for non-reductive physicalism, which, as it happens, already seemed preferrable over reductive

physicalism (e.g., Block & Fodor, 1972; Putnam, 1967). The reformulated argument reads as follows (Zhong, 2023, p. 3):

2a) [Interaction*] Mental properties are causally relevant to some physically acceptable effects.
2b) [Closure*] Every physically acceptable effect has a sufficient physically acceptable cause (at any time at which it has a cause).
2c) [Exclusion*] There is no systematic causal overdetermination by physically acceptable and physically unacceptable causes.
2d) [Identity*] Therefore, mental properties are physically acceptable properties.

Properties that are 'physically acceptable' are those properties that are either identical to or fully grounded in—and thus fully explainable in terms of—fundamental physics. The upshot is that mental properties are, at a minimum, fully grounded in physical properties.

The non-reductive physicalist thus gets exactly what she wants. The conclusion does not commit us to the identity of the mental and the physical, but it still secures that dualism, the claim that mental properties are not fully grounded in the physical, but merely necessitated by their physical correlates combined with fundamental *sui generis* psychophysical laws of nature, must be false.[1]

Zhong maintains that considerations about the nature of causation favour his reformulated argument (2a–2d) over the traditional argument (1a–1d). The traditional argument presupposes a productive or generative theory of causation (Kim, 2007, 235–236). Such accounts of causation have fallen on hard times and are typically abandoned in favour of difference-making accounts, as developed by Lewis (1973, 1979) and Woodward (2003, 2015, 2021). Zhong argues that such difference-making accounts undermine the premises of the traditional argument, while supporting the premises of the reformulated argument. I am convinced by the first part of Zhong's claim, but not the second.

The difference-making strategy against the traditional argument is quite familiar by now. According to difference-making accounts, A causes B if and only if B would not have occurred if A had not occurred: the high pollen count caused my sneeze if and only if I would not have sneezed if the pollen count was low. Zhong points out that the causal credentials of the properties studied in fundamental physics are far from impeccable on such accounts. Would I have sneezed if some of the leptons and quarks making up the pollen had been arranged slightly differently? Probably.[2] By

[1] See Chalmers (1996) and List and Stoljar (2017) for canonical formulations of such a dualist position.

[2] It is worth adding an asterisk to this intuitive response though. What fuels this intuition is a tendency to go to a scenario where the quarks and leptons are arranged slightly differently, but still in a way that gives rise to a high pollen count. Such an evaluation of counterfactuals is contentious, as Lewis and others maintain that we ought not replace the phenomena we excise with very similar phenomena (Lewis 2000; McDonnell 2019). Instead, we should consider them removed as if by a 'metaphysical hole-puncher' (Bennett 2003, p. 482). This would plausibly take out the pollen altogether, thus making it so that I do not sneeze. An important challenge for Zhong and others is to motivate a replacement reading for such counterfactuals, rather than the received hole-puncher readings.

contrast, physically grounded occurrences that are individuated at a coarser grain seem to do well for themselves as difference-makers. Would I have sneezed if the pollen count was low? Probably not.

This strategy of defanging the traditional causal argument is not only effective because difference-making accounts of causation enjoy wide support. Even if one insists that mere difference-making is not enough for causation, and some further relation of production is required, this insistence will take the sting out of causal arguments. Production theories are known to contradict common sense, and were not designed to capture our everyday judgments of what causes what (Dowe, 2000, 2004; Schaffer, 2004). But the **Interaction** premise in causal arguments is motivated mainly by such an everyday understanding of causation. The reason why this premise is taken to be non-negotiable is that mental causation plays a central role to our everyday understanding of the world, as well as much of our understanding in the special sciences (e.g., Fodor, 1989; McLaughlin, 2015; Saad, 2018). Even those defending more demanding theories of causation on theoretical grounds are happy to admit that the relation we talk and care about when we talk about causation is something much like difference-making (e.g., Dowe, 2004; Ney, 2009). So, even if, strictly speaking, mental properties fail to be causes by failing a production requirement on causation, it is unclear why a denial of **Interaction** on those grounds should rock our world.

So, as Zhong argues, difference-making accounts undermine the traditional argument. But can they support the reformulated argument? I argue that they cannot in two steps. First, I argue that difference-making accounts are too hospitable to dualist mental causation to support a causal argument against dualism. Second, I argue that Zhong's causal argument builds on a hidden premise that does not fit with difference-making accounts.

3 Real difference-making

Before taking aim at Zhong's reformulation, let me provide a broad picture of why I am sceptical of using causation as a tool for arguing against dualism.

Let us be honest about what difference-making accounts actually say about causation. They say that there is nothing to causation, over and above certain patterns of co-occurrence and co-disappearance holding. In their canonical formulations, they state that A's causing B boils down to nothing more than the fact that both A and B occurred, and if A had disappeared, then B would have disappeared as well (Lewis, 1973; Woodward, 2003). Zhong and others have argued that this condition should be supplemented with a requirement that occurrences of A should be followed by occurrences of B in relevantly similar scenarios (e.g., List & Menzies, 2009; Zhong, 2020b, a, 2022), and there are ongoing debates on how these accounts should be made more precise to deal with puzzles such as causation by absence (e.g., Henne et al., 2017; McGrath, 2005; Vaassen, 2023), redundant causation (e.g., Hitchcock, 2007, McDonald, forthcoming), and the asymmetry of causation (e.g., Albert, 2000, 2015; Lewis, 1979; Loewer, 2007). But, at bottom, all these accounts agree that there is nothing more to causation than a particular pattern of co-occurrence and

co-disappearence. These patterns may well hold because the cause and the effect are connected by some natural law, production process, or necessitation relation, but, as far as difference-making accounts are concerned, they certainly needn't be. Friends of difference-making accounts go out of their way to emphasize that one shouldn't focus on what makes it the case that A and B exhibit this pattern of correlation; as long as the pattern is out there, causation is out there (e.g., Woodward, 2021, p. 6–8).

We should thus be careful not to mistake difference-making accounts for providing merely an epistemology of causation rather than a metaphysics of causation.[3] That difference-making accounts offer a metaphysics of causation is easy to forget. Pre-theoretically, we might come to understand causation as an intrinsic process linking causes to their effects via physical contact or a transfer of energy (cf., Hall, 2004), and patterns of co-occurrence or co-disappearance can be very good indications of certain processes holding. But in order to be a philosophical position of any interest, difference-making theories must posit more than just such an evidential relation between patterns and causal relations. After all, even defenders of production accounts agree that patterns of correlation are closely related to causation (e.g., Kim, 2007). Defenders of difference-making accounts must take this one step further: these patterns are all there is to causation!

With this in mind, it is hard to see how we can raise a pressing mental causation worry for dualism. Dualism might preclude a transfer of energy from mental properties to physical properties, but, regardless of whether dualism is true, it is the case that mental properties exhibit strong patterns of correlation with effects we care about. These patterns are exactly why we are so convinced that there must be mental causation. If A causing B is just a matter of exhibiting a specific pattern of correlation, then there is no obvious problem with dualist mental causation.

Of course, it remains to be seen whether mental properties can exhibit the *right* pattern of correlation within dualist ontologies. We all know that the world is riddled with non-causal correlation patterns, such as the widespread co-occurrence between nicotine-stained finger tips and a heightened cancer risk. To settle the issue of dualist mental causation, we need to figure out whether the correlations between mental occurrences and behavioural occurrences fit the difference-making mould.

If we look at the proposed methodologies for selecting the right patterns of correlation, it seems that the dualist may rest easy. When providing a picture of how to assess 'if not A, then not B' claims such that the difference-making account picks out the right patterns of correlation, Lewis (1979) proposes to reverse-engineer the semantics of counterfactuals from our common-sense intuitions about what causes what. As champions of the **Interaction** premises will point out, the fact that our pain causes us to shriek is a 'Moorean' fact, such that is not to be given up on (cf., McLaughlin, 2015; Saad, 2018). If our commonsense intuitions are our guide, we

[3] This is of course not to say that these two projects are entirely disjoint. Woodward is adamant that our metaphysics of causation should be informed by an epistemology of causation (See Woodward 2021, Weinberger et al. 2024).

had better make the answer to 'would I have shrieked if I had not been in pain?' be 'no'. The common-sense guideline appears hospitable to dualist mental causation.

Woodward is less convinced of allotting common-sense intuitions such final authority, and proposes that usefulness in terms of reliable manipulability and prediction ought to be the mark of causal correlation patterns: causes are reliable ways to make their effects happen, and the occurrence of a cause is a reliable predictor of the occurrence of the effect (e.g.,Woodward, 2014). Again, it would seem that the dualist has nothing to fear. If we know anything about pains, it is that they are a reliable way to make people shriek and that they are reliable indicators of imminent shriekings. The usefulness guideline thus seems equally hospitable to dualist mental causation.

Of course, even if we stick to these guiding methodologies, the dualist may still be disappointed. It might turn out that, when we put these methodologies in practice, we still need to restrict the relevant patterns of correlation such that dualist mental causation is out of the question. One particular in-house dispute between difference-making accounts turns out to be of crucial importance: should the fundamental laws of nature be held fixed when we excise a purported cause to assess whether its disappearance would correlate with the disappearance of the effect? If so, the dualist is home free. If not, she is in trouble.

In the 'no' camp, we have Lewis (1979). Lewis' proposal was to excise purported causes with tiny lapses in the fundamental laws of nature (Lewis, 1979). So, when I ask, 'would I still have shrieked if I had not been in pain?', I have to look for a world that is as similar to the actual world as possible, except for the fact that my pain was excised by a tiny lapse in the psychophysical law that generated my pain based on the underlying physical property. Several authors, myself included, have pointed out that this Lewisian difference-making approach does not look promising for the dualist (Bennett, 2008; Kroedel, 2015; Vaassen, 2021b). In a world where my pain is excised by a small miracle, I would plausibly still have shrieked, because all the physical goings-on, including the physical event that necessitated my shrieking, are left intact.

Dualists should thus hope that the 'yes' camp is correct. According to this camp, fundamental laws are to be held fixed when assessing the counterfactual 'if I had not been in pain, I would not have shrieked' (e.g., Albert, 2000, 2015; Dorr, 2016; Loewer, 2007). Going by these rules, we should expect that it is true that I would not have shrieked if I had not been in pain. If we hold the fundamental laws of nature fixed, and the dualist is right that there is a fundamental psychophysical law connecting pain to physically acceptable properties, then excising my pain would require excising the physically respectable property that generated it as well. In the absence of this property (say, C-fibers firing), I would not have shrieked. If causation is difference-making, and the 'yes' camp is right that we ought to hold fundamental laws fixed when excising purported causes, then the dualist is home free with regards to mental causation.

There are plenty of reasons to believe that the 'yes' camp is in fact correct. Albert (2000, 2015), Loewer (2007), and Dorr (2016) argue for this independently of any concerns about mental causation. Dorst (2020) argues that holding laws fixed in counterfactual scenarios helps to make them relevant for reliable prediction and

control in the actual world—a much-coveted feature for difference-making accounts of the interventionist stripe. In other work, I have argued that this approach naturally delivers dualist mental causation without making faulty predictions about spurious correlations such as those between nicotine-stained fingertips and lung cancer risk, or falling barometer readings and thunderstorms (Vaassen, 2024, esp. §4.2).[4]

Rather than further weighing in on this 'yes'- 'no' debate here, I suggest we take a step back and ask what the fact that it comes down to this issue reveals about the force of causal arguments against dualism.

If the question of dualist mental causation really boils down to whether the relevant counterfactual scenarios ought to be nomologically possible or merely metaphysically possible, it is hard to see why we should care. The absence of mental causation is hailed by Fodor as 'the end of the world' (Fodor, 1989, p. 77), but think about how much of the world remains if mental causation turns out to be absent because dualism is true and Lewisian difference-making is the right account of causation. While it might not be, strictly speaking, true that my being thirsty caused me to drink, there still is a perfectly coherent sense in which it is true that if I hadn't been thirsty, I would not have drunk. It is also still true that making me thirsty is a reliable way of making me drink, and that my being thirsty is a reliable predictor of my drinking. Moreover, the perfectly coherent sense in which all these things are true matches a perfectly coherent sense in which it is true that if it weren't for the Covid pandemic, inflation would be lower and that if there had been no rain this year, my roses would not have bloomed, i.e. the sense carved out by the non-Lewisian approaches in the 'yes' camp.

The problem for devout physicalists like me is that the non-Lewisian approaches do a good job of capturing counterfactuals such as 'if the Covid pandemic had not happened, inflation would be lower', while at the same time predicting that 'if I had not been in pain, I would not have shrieked' is true regardless of whether dualism is true. This means that dualist mental properties exhibit the exact same kind of patterns of correlations with their purported effects as physically respectable properties exhibit with their purported effects. Whether or not that exact pattern of correlations rather than the pattern picked out by the Lewisian approach deserves the label 'causation' seems to be of very little interest. Both patterns support reliable prediction and control in the world. They both quack and walk like causation to such an extent that it is hard to see why we should care which one is the real deal.

If the traditional causal argument ran on productive accounts of causation, and the revised causal argument runs on Lewisian difference-making, the upshot of both is bound to be painfully similar. Even if they manage to establish that dualism is incompatible with mental causation (which I doubt they do), it is hard to see how it burdens the dualist with a real problem. It's not the end of the world, it's a slight mistake in our labelling.[5]

[4] See also Kroedel (2015, 2020), who argues that even the 'no' camp can provide a safe haven for the dualist, if she is willing to give psychophysical laws a special dispensation.

[5] Similar concerns with causal arguments are voiced by Chalmers (1996) and Pautz (forthcoming, §8).

4 Zhong's argument

If I am right that causation is the wrong tool to whittle down dualism to an untenable position, there must be something wrong with causal arguments of the kind that Zhong proposes. What, then, goes wrong in Zhong's causal argument?

The problem with Zhong's argument is somewhat hidden. The dualist can agree with **Interaction***, **Closure***,[6] and **Exclusion***.[7] The real problem lies in a hidden assumption about what overdetermination consists of.

Zhong's argument assumes that if there is a physically respectable event PC that is a sufficient cause of the target effect E, then any non-physical event that causes E and occurs simultaneously with PC must overdetermine E. This assumption is hard to motivate, and easily rejected by the dualist. Suppose that C-fibers firing is a suitably sufficient physically acceptable cause of my shrieking. Suppose also that dualism is true and my pain is not fully grounded by C-fibers firing, but naturally necessitated by it in conjunction with a primitive psychophysical law. Now, is it true that if both my pain and my C-fibers firing cause me to shriek, my shrieking is overdetermined? Not if we take 'overdetermination' to mean an objectionably *ad hoc* case where two independent causal chains bring about the same effect, thereby making one another redundant. We need no *ad hoc* extra mechanism connecting the pain to the shrieking: the relevant pattern of correlation already holds in virtue of my pain being connected to my C-fibers firing, and my C-fibers being connected to my shriek-muscles. The case is clearly different from one where two distinct rocks simultaneously shatter the same window.[8] If we want to call this a case of overdetermination, it should be qualified as a benign case that poses no extra burden on our ontology (cf., Bennett, 2003; Pereboom, 2002; Sider, 2003).

The following picture arises. There are physically respectable, suitably sufficient causes for all effects, as per **Closure***. There are also non-physical causes for behavioural effects, as per the combination of dualism and **Interaction***. There is no systematic problematic overdetermination, as per **Exclusion***. There is, however, a benign overdetermination of behavioural effects: they are caused by both physically respectable, suitably sufficient causes such as C-fibers firing and simultaneously occurring non-physical mental events such as pain.

[6] Though I think she need not accept **Closure***. There is evidence that physically respectable but non-fundamental events will always fail to be sufficient for their target effects. Any mid-sized event, such as my throwing a rock, can be physically realized such that it has a thermodynamically abnormal future (cf., Albert, 2000, 2015; Field, 2003). For example, my rock throw can be realized such that the rock will change into a sponge mid-trajectory, thus leaving the window unshattered—if lightly smudged. So, as it stands, **Closure*** is probably false. It won't do the dualist much good to object in this way however, as there is no reason to believe that mental phenomena are waiting in the ranks to close the causal gaps left by higher-level phenomena. I have argued elsewhere that this problem with causal arguments can be patched without creating further trouble for physicalists (Vaassen, 2021a).

[7] See Zhang (2024) for a further argument that the denial of dualism does not logically follow from these three premises alone.

[8] See Kroedel (2015, §5), Vaassen (2024, §4.3), and (Rubenstein, forthcoming, ms) for a more precise account of how they differ.

Zhong provides two reasons for objecting to such a picture: one might be sceptical of the possibility of such benign overdetermination across the board and, even if there is benign overdetermination, one might be sceptical that dualist mental properties and their underlying physical properties are good candidates for co-causing in such a benign way. I'll take these points in reverse order.

First, Zhong argues that dualist mental properties and their underlying physically respectable properties cannot co-cause in the way of other cases of benign overdetermination, because there is no full grounding relation between them. Given that non-physical properties are not fully grounded in physically respectable properties, the former cannot inherit the causal powers of the latter. Similarly, he maintains that the causal powers of the physically acceptable properties cannot explain the causal powers of non-physical properties without a full grounding relation between them. All that the nomic relations that the dualist allows between physically acceptable and non-physical properties can do is explain why they systematically co-occur (Zhong, 2023, p. 5–6). However, if we buy into difference-making accounts, causal inheritance is an inheritance of co-occurrence and co-disappearance patterns. Consequently, explaining causal powers is just a matter of explaining co-occurrence and co-disappearance patterns. The correlation patterns of mental phenomena, and how they are inherited from the physical properties underlying them, are exactly the kind of facts that we should expect to explain their causal status.

Second, Zhong has his reservations about benign overdetermination across the board. Considering cases such as my pain and its underlying physical property, Zhong states that the pain-shrieking correlation must be either sensitive, such that I would not have shrieked if the pain was due to a slightly different physical property, or insensitive, such that I would have shrieked if the pain was due to a slightly different physical property. If the pain-shrieking correlation is sensitive, then the underlying physical property is the cause of my shrieking, and my pain is not. If the pain-shrieking correlation is insensitive, then my pain is a cause of my shrieking, and the underlying physical property is not. Either way, we avoid having two causes at once, and we thus avoid overdetermination of any kind.

I think there is reason to believe that Zhong's one-cause-only policy is too restrictive. Socrates died by drinking poison. He also died by drinking hemlock, as that was the poison in question (cf., Bontly, 2005; McDonnell, 2017). There is nothing awfully *ad hoc* or burdensome about accepting that both the drink being poison and its being hemlock caused him to die. Similarly, I see no reason why it is *ad hoc* or burdensome to accept that both my pain and my C-fibers firing cause me to shriek. In cases where the correlation between top-level property and the target effect is insensitive, benign overdetermination seems perfectly acceptable to me (cf., Vaassen, 2022).

If, despite this, the dualist wants to avoid even benign overdetermination at all costs, she can imitate Zhong's own strategy for doing so. If the correlation between my pain and my shrieking is sensitive to changes in my C-fibers firing, then my C-fibers are a cause and my pain is not. If the correlation between my pain and my shrieking is insensitive, then my pain is a cause and my C-fibers are not. Of course, this might mean that **Closure*** is false, but note that it would be false in the same innocent way that **Closure** is false: the events actually sufficient for the behavioural

effect are not denied their status as being sufficient, they are being denied the label 'cause'. If we are willing to deny physical events at the fundamental level that status when there are non-fundamental events in the neighbourhood that are better difference-makers, I don't see why we should be reluctant to deny physically respectable events that status when there are non-physical events in the neighbourhood that are better difference-makers. Because I myself see no reason to be coy with the label 'cause', this would not be my personally preferred strategy if I were hit over the head and became a dualist, but it is available for those who prefer draconian labelling practices.

Of course, the picture I have just defended presupposes that we can allow for dualist mental causation via the use of non-Lewisian difference-making as outlined above. If I am mistaken about this, the dualist must give up on **Closure***, **Interaction***, or **Exclusion***. However, this outcome would not help us get right with God about our physicalist commitments. Now that we are wise to the superficiality of causation, we know that giving up on **Interaction*** is not the humiliating defeat it was once taken to be. As prefigured in §3, difference-making accounts take the sting out of causal arguments, even those that target dualism.

5 Conclusion

If difference-making accounts of causation are correct, causation is a superficial relationship. As such, it is a poor guide to the metaphysically weighty facts of the world, such as what is identical to what, and what metaphysically necessitates what. I submit that this finding undermines Zhong's causal argument for physicalism in philosophy of mind. Note, however, that the concerns I raised here are likely to affect other debates as well. In metaethics and aesthetics, for example, the causal efficacy of properties such as goodness and beauty is taken to stand or fall with their being tightly tethered to physical or natural properties (Enoch, 2017; Thomson, 1996, Evers, 2024). In as far as the underlying reasoning mimics the causal arguments considered here, we should expect this assumption to be mistaken. In these debates, as in philosophy of mind, a re-evaluation of what the causal facts can actually tell us about the metaphysical facts is long overdue.[9]

Funding Open access funding provided by Umea University.

Data availability My manuscript has no associated data.

Declarations

Ethics approval No specific approvals of ethical boards were required for this research.

[9] I am grateful to the editorial team of this journal for inviting me to contribute to this special issue, as well as to an anonymous referee who delivered helpful suggestions for improving the text. Further, I thank Karen Bennett and Audrey Powers for their comments on earlier versions of this article.

Conflict of interest The author declares no competing interests.

References

Albert, D. Z. (2000). *Time and chance*. Harvard University Press.
Albert, D. Z. (2015). *After physics*. Harvard University Press.
Bennett, K. (2003). Why the exclusion problem seems intractable, and how, just maybe, to tract it. *Noûs, 37*(3), 471–497.
Bennett, K. (2008). Exclusion again. In J. Kallestrup & J. Hohwy (Eds.), *Being reduced* (pp. 280–307). Oxford University Press.
Block, N. (2003). Do causal powers drain away? *Philosophy and Phenomenological Research, 67*(1), 133–150.
Block, N., & Fodor, J. A. (1972). What psychological states are not. *Philosophical Review, 81*(April), 159–181.
Bontly, T. D. (2005). Proportionality, causation, and exclusion. *Philosophia, 32*(1), 331–348.
Chalmers, D. J. (1996). *The conscious mind*. Oxford University Press.
Dorr, C. (2016). Against counterfactual miracles. *Philosophical Review, 125*(2), 241–286.
Dorst, C. (2020). Why do the laws support counterfactuals? *Erkenntnis, 87*(2), 545–566.
Dowe, P. (2000). *Physical causation*. Cambridge University Press.
Dowe, P. (2004). Causation and misconnections. *Philosophy of Science, 71*(5), 926–931.
Enoch, D. (2017). Non-naturalistic realism in metaethics. In T. McPherson & D. Plunkett (Eds.), *The Routledge handbook of metaethics* (pp. 29–42). Routledge.
Evers, D. (2024). Aesthetic Non-Naturalism. *The British Journal of Aesthetics*. https://doi.org/10.1093/aesthj/ayad047
Field, H. (2003). Causation in a physical world. In M. J. Loux & D. W. Zimmerman (Eds.), *The Oxford handbook of metaphysics* (pp. 435–460). Oxford University Press.
Fodor, J. A. (1989). Making mind matter more. *A theory of content and other essays* (pp. 137–159). MIT Press.
Gebharter, A. (2017). Causal exclusion and causal bayes nets. *Philosophy and Phenomenological Research, 95*(2), 353–375.
Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, & L. Paul (Eds.), *Causation and counterfactuals* (pp. 225–276). MIT Press.
Henne, P., Pinillos, A., & Brigard, F. D. (2017). Cause by omission and norm: Not watering plants. *Australasian Journal of Philosophy, 95*(2), 270–283.
Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *Philosophical Review, 116*(4), 495–532.
Kim, J. (1989). The myth of non-reductive materialism. *Proceedings and Addresses of the American Philosophical Association, 63*(3), 31–47.
Kim, J. (2005). *Physicalism, or something near enough*. Princeton University Press.
Kim, J. (2007). Causation and mental causation. In B. P. McLaughlin & J. D. Cohen (Eds.), *Contemporary debates in philosophy of mind* (pp. 227–242). Blackwell.
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. (1973). Causation. *The Journal of Philosophy, 70*(17), 556–567.
Lewis, D. (1979). Counterfactual dependence and time's arrow. *Noûs, 13*(4), 455–476.

Lewis, D. (2000). Causation as influence. *Journal of Philosophy, 97*(4), 182–197.
List, C., & Menzies, P. (2009). Non-reductive physicalism and the limits of the exclusion principle. *The Journal of Philosophy, 106*(9), 475–502.
List, C., & Stoljar, D. (2017). Does the exclusion argument put any pressure on dualism? *Australasian Journal of Philosophy, 95*(1), 96–108.
Loewer, B. M. (2007). Counterfactuals and the second law. In H. Price & R. Corry (Eds.), *Causation, physics, and the constitution of reality: Russell's republic revisited* (pp. 293–326). Oxford University Press.
McDonald, J. (forthcoming). Essential structure for causal models. *Australasian Journal of Philosophy.*
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, S. (2005). Causation by omission: A dilemma. *Philosophical Studies, 123*(1–2), 125–148.
McLaughlin, B. P. (2015). Does mental causation require psychophysical identities? In T. Horgan, M. Sabates, & D. Sosa (Eds.), *Qualia and mental causation in a physical world: Essays in honor of Jaegwon Kim.* Cambridge University Press.
Ney, A. (2009). Physical causation and difference-making. *British Journal for the Philosophy of Science, 60*(4), 737–764.
Papineau, D. (2001). The rise of physicalism. In C. Gillett & B. M. Loewer (Eds.), *Physicalism and its discontents* (pp. 3–36). Cambridge University Press.
Papineau, D. (2002). *Thinking about consciousness*. Oxford University Press.
Pautz, A. (forthcoming). How to achieve the physicalist dream: Identity or grounding? In Rabin, G., (Ed.), *Grounding and consciousness.* Oxford University Press
Pereboom, D. (2002). Robust non-reductive materialism. *Journal of Philosophy, 99*(10), 499–531.
Putnam, H. (1967). The nature of mental states. In W. Capitan & D. Merrill (Eds.), *Art, mind, and religion* (pp. 1–223). Pittsburgh University Press.
Rubenstein, E. (forthcoming). Avoiding overdetermination. https://tinyurl.com/Avoiding-Overdetermination. Accessed 13 May 2023.
Russo, A. (2016). Kim's dilemma: Why mental causation is not productive. *Synthese, 193*(7), 2185–2203.
Saad, B. (2018). A causal argument for dualism. *Philosophical Studies, 175*(10), 2475–2506.
Schaffer, J. (2004). Causes need not be physically connected to their effects: The case for negative causation. In C. R. Hitchcock (Ed.), *Contemporary debates in philosophy of science* (pp. 197–216). Basil Blackwell.
Sider, T. (2003). Review: What's so bad about overdetermination? *Philosophy and Phenomenological Research, 67*(3), 719–726.
Thomson, J. J. (1996). Epistemological arguments for moral scepticism. In G. Harman & J. J. Thomson (Eds.), *Moral relativism and moral objectivity.* Blackwell.
Vaassen, B. (2021a). Causal exclusion without causal sufficiency. *Synthese, 198*, 10341–10353.
Vaassen, B. (2021b). Dualism and exclusion. *Erkenntnis, 86*(3), 543–552.
Vaassen, B. (2023). Absence and abnormality. *Analysis, 83*(1), 98–106.
Vaassen, B. (2022). Halfway proportionality. *Philosophical Studies*, *179*, 2823–2843.
Vaassen, B. (2024). Mental causation for standard dualists. *Australasian Journal of Philosophy*, 1–21. https://doi.org/10.1080/00048402.2024.2335325
Weinberger, N., Williams, P., & Woodward, J. (2024). The worldly infrastructure of causation. *British Journal for the Philosophy of Science.* https://doi.org/10.1086/730698
Woodward, J. (2003). *Making things happen: A theory of causal explanation*. 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), 691713.
Woodward, J. (2015). Interventionism and causal exclusion. *Philosophy and Phenomenological Research, 91*(2), 303–347.
Woodward, J. (2021). *Causation with a human face: Normative theory and descriptive psychology*. Oxford University Press.
Zhang, W. (2024). On the causal arguments for physicalism. *Asian Journal of Philosophy, 3*(1), 1–11.
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.
Zhong, L. (2022). Why causation is biconditional but not proportional. *American Philosophical Quarterly, 59*(3), 263–271.
Zhong, L. (2023). A causal argument for physicalism. *Asian Journal of Philosophy, 2*(2), 1–9.