

Overdetermination and Causal Connections

Some theories are alleged to be implausible because they are committed to systematic ‘overdetermination’. In response, some authors defend ‘compatibilism’: the view that the putative overdetermination is benign, like other unproblematic cases of a single effect having many sufficient causes.

The literature (e.g. Yablo 1992, Kim 1993, Bennett 2003, Wilson 2009, and List & Stoljar 2017) has tended to focus on the following question: which relations between sufficient causes of a single effect ensure that problematic overdetermination is avoided? This paper argues that several widely endorsed answers to this question are subject to counterexample. It then proposes a diagnosis of this failure: the standard answers neglect what really matters — how the causes are connected to their shared effect. In particular, overdetermination is avoided when there are no independent causal connections.

According to the ‘Busy God’ hypothesis, all effects have both a natural cause and a divine cause.¹ When the window shatters, it is not only because the rock hit it, but also because God willed it to shatter. When the match lights, it is not only because it was struck, but also because God willed it to light. And so on. The Busy God hypothesis is extravagant — in a bad way. A familiar diagnosis of this extravagance is that it involves systematic *overdetermination*. It systematically posits two causes where only one is needed; it multiplies causes beyond necessity.

Systematic overdetermination counts especially heavily against a theory, but even a single case of overdetermination can be damaging. Holmes and Watson are puzzling over Victim’s death: Holmes envisages a single sniper, whereas Watson posits two, operating independently, who happened to strike at precisely the same moment. Watson’s theory makes Victim’s death overdetermined. Thus, Holmes’s theory is *prima facie* preferable; Watson’s cannot be ruled out, but it requires strong evidence to be plausible. Perhaps, in rare cases, murders involve overdetermination of the sort that Watson describes, but we don’t expect many to. Positing overdetermination is costly and inelegant; positing it systematically makes a theory incredible.

¹ This example is from Bennett (2003: 475).

This raises the following question: when is a theory committed to overdetermination of this implausible kind? The literature offers several widely endorsed ‘overdetermination-avoiding conditions’, according to which distinct causes do not overdetermine their shared effect so long as they are suitably related. In particular, causes C1 and C2 do not overdetermine an effect E if C1 causes C2, or if C1 grounds C2, or if C1 requires C2 for E, in the sense that E would not have occurred if C1 had occurred without C2. After some setup (§1), this paper argues that these conditions are each subject to counterexample (§2). It then proposes a diagnosis (§3): the standard answers neglect what really matters — how the causes are connected to their shared effect. In particular, overdetermination is avoided when there are no independent causal connections.

1. Preliminaries

Since it is plausible that any physical effect has a sufficient microphysical cause, any theory which posits high-level causes of physical effects is in principle open to the charge of systematic overdetermination. ‘Causal exclusion’ arguments exploit this idea. They have been especially prominent in the philosophy of mind, where they are the catalyst for the rise of ‘reductive’ physicalism (Papineau 2001), the most serious problem for dualism (Chalmers 2011), and the source of debate over the intermediate position of ‘non-reductive’ physicalism (Kim 1993, 1998; Bennett 2003, 2008).² Mental events (e.g. pains) and physical events (e.g. neural firings) seem to systematically share common effects (e.g. wincing). Thus, mental events and physical events surely have *something* to do with each other; otherwise, our theory of mental causation is like the Busy God hypothesis, or Watson’s double-sniper story writ large.

I use the expression ‘C is a sufficient cause of E’ (equivalently, ‘C is causally sufficient for E’), to mean that C is a cause of E which, together with the relevant background conditions and causal laws, guarantees E’s occurrence. For example, the sniper’s shooting is a sufficient cause for Victim’s death, even though strictly speaking the relevant causal chain requires many background

² As I use this terminology, reductive physicalism holds that mental events are physical events, whereas non-reductive physicalism holds that they are distinct from but metaphysically dependent on physical events.

conditions to be in place: Victim must not be wearing a bullet-proof vest, or have a bodyguard disposed to jump into the bullet's path, and so on.³

Call any case in which an event has more than one sufficient cause 'multi-determination'. The exclusion arguments I focus on instantiate the following simplified schema:

1. Theory T is committed to systematic multi-determination of kind K.

2. Systematic multi-determination of kind K is implausible.

Therefore, T is implausible.

For example, where T is grounding physicalism, the multi-determination in question involves mental and physical causes for behavior, where the physical cause grounds the mental cause. Where T is dualism, the multi-determination in question involves mental and physical causes which are metaphysically distinct (though they may be nomically connected). In each case, the question is whether the theory is systematically committed to the problematic kind of multi-determination — the kind which is exhibited by the Busy God hypothesis and Watson's double-sniper story. I will use 'overdetermination' to denote this kind of multi-determination: in this terminology, the objection is that the theory is committed to systematic overdetermination.⁴

One kind of response rejects the first premise, by denying that the alleged overdetermination is even multi-determination. Perhaps physics is causally incomplete (so that the physical effects of mental/high-level events lack microphysical causes), or perhaps mental/high-level events are 'epiphenomenal' (so that the physical effects of micro-physical events lack mental/high-level

³ There is a more demanding notion of causal sufficiency, on which sufficient causes physically necessitate their effects. This is more demanding both because it involves physical laws, rather than higher-level 'causal laws' (which may be physically contingent), and because it does away with background conditions. This is not the notion that is relevant to typical exclusion arguments, since ordinary high-level events do not physically necessitate anything other than highly complex disjunctions of microphysical conditions. (See Field 2003: 439 and Albert 2015: 2.)

⁴ 'Overdetermination' is sometimes used to mean (simultaneous) multi-determination (or some variant of it), with the crucial distinction between 'benign' overdetermination and 'bad' overdetermination. For what it's worth, this terminology seems distorted: oversleeping, overspending, and overcooking are all inherently excessive.

causes). This approach strikes me, at least in typical cases, as a last resort: it seems overwhelmingly plausible both that physics is causally complete and that higher-level events have physical effects.⁵

My interest is therefore in an alternative response: ‘compatibilism’ (Bennett 2003: 473). Compatibilists deny the second premise: they hold that the multi-determination in question is not *overdetermination* but is instead ‘benign’. After all, multi-determination is ubiquitous: effects typically have very many sufficient causes. Thus, compatibilists aim to identify sufficient conditions for multi-determination to be benign — as I will call them, ‘overdetermination-avoiding conditions’ (OACs) — and to argue that these conditions are met by cases of multi-determination of kind K.⁶

A crucial question for compatibilists is: which conditions are OACs? To simplify, I follow standard practice by focusing on conditions of the following form:⁷

C1 and C2 do not overdetermine E if

⁵ Some posit ‘downwards exclusion’, according to which physics fails to be causally complete since microphysical events are too specific to count as causes (e.g. List & Menzies 2009, Zhong 2020). However, as Vaassen (2021) points out, this doesn’t address the exclusion problem. Let a ‘sufficer’ for E be an event which, together with the relevant background conditions and causal laws, guarantees E’s occurrence, and which satisfies some minimal condition for being part of E’s causal history (e.g. by grounding a cause of E). Call any case in which an event has more than one sufficer ‘multi-determination*’. Then downwards exclusion views still posit systematic multi-determination*, and so they face the analogous question of when multi-determination* is ‘benign’.

The implicit idea behind these views may be that problematic overdetermination is avoided when only one of the sufficers is a (proportional) cause. Of course, the action will then be in spelling out the relevant notion of ‘cause’. The discussion below can be read as suggesting that the question of whether two events are each causes (in this sense) may turn at least in part on whether their causal connections to the effect are suitably independent.

⁶ E.g. Yablo 1992, Shoemaker 2001, Pereboom 2002, Bennett 2003, Wilson 2009, Crane & Árnadóttir 2013, Kroedel & Schulz 2016, List & Stoljar 2017.

⁷ This is not to say that the generalizing from this case is trivial: whether some sufficient causes C1, ..., Cn overdetermine E may not merely be a matter of whether they pairwise overdetermine E.

Two observations inspire the search for OACs. First, in paradigm cases of benign multi-determination, one sufficient cause is itself a cause of the other. For example, the rock-throwing and the rock-striking multi-determine the window-shattering. Second, one relation between sufficient causes which obviously avoids overdetermination is identity: if the physical event and the mental event turn out to be one and the same, then there is no *multi*-determination of behavior, and *a fortiori* no *overdetermination*. This motivates the idea that multi-determination is benign so long as the causes stand in a sufficiently intimate identity-like relation. In particular, many have thought that overdetermination is avoided if high-level events are not ‘metaphysically distinct’ from underlying microphysical events, perhaps by supervening on or being grounded in them (cf. Papineau 2002: 33; List & Stoljar 2017).

Some discussions of causal exclusion focus on *impossibility* (as a matter of metaphysical principle): the worry is that it might somehow follow from the nature of causation itself that sufficient micro-physical causes render high-level events causally irrelevant. This is suggested by talk of there being ‘no causal work left over’ for high-level events (Kim 1993: 354), of their ‘causal powers draining away’ to lower levels (Kim 1998: 81), or of events ‘competing’ for causal influence (Yablo 1992: 272). Many authors have offered convincing responses to this kind of concern. Several have argued that counterfactual-based approaches to causation — especially as formulated within the interventionist framework — show that it is perfectly possible for high-level events to be causally relevant despite the sufficiency of their low-level realizers.⁸

It is obviously a different task to show that a given kind of multi-determination is *plausible* (as a matter of epistemic practice). For example, we might grant dualists that the nature of causation allows in principle that, whenever some brain state brings about some behavior, that behavior is also brought about by some accompanying mental state. But we might think that this theory is no better than the Busy God hypothesis.⁹ This, therefore, is the question I will focus on: when does a

⁸ See, for example, Lepore & Loewer 1987, Shapiro & Sober 2007, Woodward 2008, 2015, List & Menzies 2009, Shapiro 2012, Rescorla 2014, Weslake forthcoming. As Sider 2003 points out, all standard accounts of causation allow for the possibility of double-sniper cases.

⁹ A subtle slide from the possibility-based form of exclusion argument to the plausibility-based form is achieved by supplementing an exclusion principle according to which multi-determination of kind K is impossible with the crucial caveat: ‘unless it is a case of genuine overdetermination’ (Kim 2005: 42), where

case of multi-determination exemplify the kind of problematically implausible overdetermination, as exhibited paradigmatically by the double-sniper and Busy God cases?

2. Counterexamples

This section presents counterexamples to three natural and commonly endorsed OACs: cases which satisfy the condition in question and yet, intuitively, involve overdetermination (in the problematic sense introduced above). When evaluating these cases, it is important to remember what I mean by the claim that a case involves ‘overdetermination’. I do not mean that the case is impossible, or that we couldn’t be justified in believing that it happened. I simply mean that it involves the kind of multi-determination which could not plausibly be posited systematically. (In fact, I mainly rely on variations of the double-sniper case, for which the question is: are these variations relevantly different from the original?)

2.1 Causing

C1 is causally sufficient for E; so is C2. But no matter: C2 is itself an effect of C1. So go paradigm cases of benign multi-determination, like:

(BILLIARD BALLS) White hits Red. Red hits Black. Black moves.

White’s collision with Red and Red’s collision with Black are each sufficient causes of Black’s motion; but since the first collision causes the second, there’s nothing excessive here. Such cases suggest the following OAC:¹⁰

Causing: C1 and C2 do not overdetermine E if C1 causes C2.

‘genuine overdetermination’ here means a case which is relevantly similar to the double-sniper case or the Busy God hypothesis (and so implausible, especially when replicated systematically).

¹⁰ Cf. Goldman (1969: 471-2), Yablo (1992: 272).

But consider the following version of the double-sniper case:

(STARTLED SNIPER) Sniper 1 and Sniper 2 each have their finger on the trigger, ready to shoot. Sniper 1 shoots; Sniper 2, startled by the sound of the shot, pulls her trigger too...¹¹

Thus, Sniper 1's shot causes Sniper 2's shot. But intuitively, it remains an extravagant coincidence that their bullets arrive at Victim simultaneously: Victim's death remains overdetermined.¹²

You may be thinking: STARTLED SNIPER only involves overdetermination because there are pairs of *intermediary* events (e.g. Sniper 1's bullet striking Victim and Sniper 2's bullet striking Victim) which overdetermine Victim's death. Hence, Victim's death is not strictly speaking overdetermined by the firings (which are appropriately related, after all) but only by these downstream intermediaries (which are not). Since this response is also relevant to the other two counterexamples discussed in this section, I postpone my argument against it to §2.4.

2.2 Grounding

C1 causes E; so does C2. But no matter: C2 is grounded in C1. So go paradigm cases of benign multi-determination, like Stephen Yablo's (1992) familiar:

(SOPHIE) Sophie the pigeon pecks whenever a red stimulus is presented. A scarlet stimulus is presented. She pecks.

¹¹ Bennett (2003: 478) uses this kind of counterexample to motivate her counterfactual criterion (discussed below).

¹² Nor does C1 and C2's sharing a common cause suffice to avoid overdetermination. Suppose that Sniper 1 and Sniper 2 each pull the trigger as a result of being startled by a thunderclap; intuitively, Victim's death remains overdetermined.

The stimulus's being scarlet and the stimulus's being red are each sufficient causes of Sophie's pecking; but since the first grounds the second, there is nothing excessive here.¹³ (Under a coarse-grained conception of events, the stimulus's being scarlet is identical to the stimulus's being red. But for the sake of argument I assume a suitably fine-grained conception: we are considering the compatibilist position that posits genuine multi-determination but no overdetermination.)

Such cases suggest the following OAC:¹⁴

Grounding: C1 and C2 do not overdetermine E if C1 grounds C2.

But consider the following double-sniper case:

(RED FLAG) Sniper 1 is under instructions to shoot if she sees a scarlet flag; Sniper 2 is under instructions to shoot if she sees a red flag. A scarlet flag is waved...¹⁵

The waving of the scarlet flag is causally sufficient for Victim's death, and so is the waving of the red flag. But the latter event's being grounded in the former event does not seem to remove the implausible coincidence that they lead to distinct bullets striking Victim simultaneously. This is not the sort of causal structure we expect to be systematically replicated!¹⁶

¹³ Some, following Yablo himself, have thought that there is not even multi-determination in such cases, since the grounding event is too specific to count as a cause. As explained in n.5 above, however, this only makes a terminological difference: cases like SOPHIE satisfy a less demanding notion of 'multi-determination*', and we can ask when cases satisfying this notion count as overdetermination.

¹⁴ Kroedel & Schulz (2016). Yablo (1992), Shoemaker (2001), Pereboom (2002), Bennett (2003), Wilson (2009) and Crane & Árnadóttir (2013) defend similar positions, with respect to determination, realization, constitution, and metaphysical necessitation.

¹⁵ Cf. Bernstein's (2016: 36) 'Alarm' case.

¹⁶ Nor does C1 and C2's sharing a common ground suffice to avoid overdetermination. Sniper 1 is under instructions to shoot if she sees a flag with an even number on it; Sniper 2 if she sees a flag with a prime number on it. A flag with '2' on it is waved: this event grounds both the waving of an even-numbered flag and the waving of a prime-numbered flag. But intuitively, these events overdetermine Victim's death.

Schaffer 2015 (amongst others) argues that grounded events do not count against the parsimony of a theory. This might seem to support *Grounding*, since the implausibility of systematic overdetermination rests, at least in part, on its lack of parsimony. However, the cost of the grounded event's occurring is one thing; the cost of its causing the same effect as the event which grounds it is another. In RED FLAG, what is surprising is not that there is a red flag-waving in addition to the scarlet flag-waving, but that both events lead, via independent routes, to Victim's death. (Analogously, suppose that grounding physicalists elaborate their theory by claiming that brain states cause behavior via one nervous pathway, and the mental states which they ground cause behavior via another independent nervous pathway. On such a theory, behavior seems implausibly overdetermined.)¹⁷

Again, you may be thinking: RED FLAG only involves overdetermination by intermediaries (e.g. the two bullet-strikings) which are not related by grounding. Before arguing against this response, let me discuss a final proposal.

2.3 Requiring

A natural and popular idea is to characterise OACs in terms of counterfactuals. We already saw in STARTLED SNIPER that it won't be enough for one sufficient cause to counterfactually depend on the other. But we might look instead towards counterfactuals involving the effect.

The idea is that overdetermination involves a kind of redundancy. In extreme cases, one of the sufficient causes makes no difference whatsoever to the effect: if it were removed, whilst holding the other sufficient cause fixed, the effect would remain the same in all details. (This might seem to be the case with dualist mental causation, for example: remove the mental cause whilst leaving the physical cause intact and no change in behavior results.) But avoiding extreme redundancy

¹⁷ Kroedel & Schulz (2016: 1914) hold that mental events cause their physical effects *because* the physical events which ground them cause these effects (though this plays no official role in their response to exclusion). As I argue below, something like this condition is what allows overdetermination to be avoided in such cases: grounded causes typically do not overdetermine because there is typically also a grounding relation between the 'causal connections' involved.

isn't sufficient for avoiding overdetermination: each sniper's firing makes a difference to the details of Victim's death.¹⁸

Double-sniper cases do exhibit a more moderate form of redundancy, however: remove either shot, and Victim's death would remain (though not in all of its details). In such cases, we can say that neither cause 'requires' the other for the effect — that is, the following counterfactuals are both false:

- i. $C1 \ \& \ \sim C2 \ > \ \sim E$
- ii. $\sim C1 \ \& \ C2 \ > \ \sim E$

Thus, we might propose that avoiding not just extreme but also moderate redundancy suffices to avoid overdetermination — in OAC-form:¹⁹

Requiring: C1 and C2 do not overdetermine E if C1 requires C2 for E

i.e. if $C1 \ \& \ \sim C2 \ > \ \sim E$.

¹⁸ Nor is it necessary, at least as it stands: proximal causes (like Red's hitting Black) screen off distal causes (like White's hitting Red), so that removing the distal cause whilst leaving the proximal cause intact makes no difference whatsoever to the effect.

¹⁹ This has been influentially defended by Karen Bennett (2003, 2008). As Bennett observes (2003: n.8), this OAC is widely accepted (albeit, usually implicitly): she cites, amongst others, Lewis (1973: 193), Lepore & Loewer (1987: 639), and Mills (1996: 107). (Lewis's definition of 'overdetermination' implies the stronger claim that C1 and C2 would each *cause* E in the absence of the other.) It is important to note, however, that these authors may not all be using 'overdetermination' in the pejorative sense at issue here. (Indeed, Schaffer (2003: 27–9) argues that Lewisian 'overdetermination' is ubiquitous, and should therefore be embraced by compatibilists.)

Technically, Bennett defends the stronger claim that overdetermination is avoided if ' $C1 \ \& \ \sim C2 \ > \ E$ ' is false. Unlike the OAC in the text, this is susceptible to counterexamples based on failures of counterfactual excluded middle. Two assassins independently plant imperfectly reliable bombs which go off at the same time, killing Victim. The counterfactual 'had only Assassin 1 planted a bomb, Victim would still have died' is false, since the bomb (being imperfectly reliable) might not have gone off. Yet Victim's death is overdetermined.

This successfully accommodates cases of benign multi-determination like the following:²⁰

(PHOTORECEPTOR) When light within a certain frequency range hits a photoreceptor cell, it produces a current.

The frequency of the incoming light is causally sufficient for the photoreceptor cell producing a current. The light's wavelength is also causally sufficient: holding fixed the refractive index of the surrounding medium, its wavelength determines its frequency. Intuitively, however, the frequency event and the wavelength event do not overdetermine the current. And indeed, the case passes the counterfactual test: if the refractive index of the surrounding medium were altered, and the incoming light had fallen outside the frequency range but not the wavelength range, then the photoreceptor would not have produced a current. Hence, the wavelength event requires the frequency event for the current. This case is not covered by *Grounding*, since the wavelength and frequency events are related nomically, not metaphysically. Nor is it covered by *Causing*, since the events are simultaneous (and, arguably, not fully distinct).

Such cases suggest that one sufficient cause requiring another makes their multi-determination of a common effect benign. But consider the following (adapted from Aimar 2012):²¹

(NEAR MISS) Just as Sniper 1 shoots, a drone happens to be flying past, en route to collide with her bullet if left unchecked. But Sniper 2's shot startles the drone's

²⁰ It also handles cases like BILLIARD BALLS (the distal cause requires the proximal cause). Its verdict on grounding cases like SOPHIE is unclear, since one of the counterfactuals has an impossible antecedent. Bennett (2003) holds that overdetermination is also avoided in these cases: her overall view is that there is only overdetermination if the counterfactual ' $\sim C1 \ \& \ C2 > E$ ' is *non-vacuously* true, and hence is avoided when this counterfactual is either false or only vacuously true. Since this OAC faces the same counterexamples as *Grounding*, I focus on non-vacuous cases here.

²¹ Aimar's case involves a protective shield being deactivated by Sniper 1's shot — I have replaced the shield with a passing drone to make it clearer that Sniper 2's shot is causally sufficient. Won (2014: 214) independently describes a similar counterexample. It is also somewhat similar to a case discussed by Strevens (2007: 101), though his involves pre-emption rather than overdetermination.

remote controller who is standing nearby – the drone swerves, avoiding Sniper 1’s bullet. Both bullets strike Victim simultaneously.

Sniper 1’s firing requires Sniper 2’s firing for Victim’s death: if Sniper 1 had fired without Sniper 2 firing, the drone wouldn’t have swerved, and Sniper 1’s bullet would have collided with it, saving Victim.²² We shouldn’t say that Sniper 1’s shot is not causally sufficient for Victim’s death, however: it guarantees it in the absence of any miraculously intervening drones — surely a reasonable background condition to include when assessing causal sufficiency. Thus, there is multi-determination, but overdetermination is avoided by the lights of *Requiring*. Intuitively, this is the wrong result: it remains an extravagant coincidence that Victim’s death was caused by Sniper 1’s firing and Sniper 2’s firing.²³ (It would be very surprising if mental causation worked like *that!*)²⁴

Once again, you may be thinking: NEAR MISS only involves overdetermination because of downstream intermediaries.²⁵ Now is the time to address this response.

2.4 The intermediaries strategy

A natural response to each of the counterexamples above accepts the verdict of the OAC in question that the initiating causes are not themselves overdeterminers, and blames the

²² This is what Hall (2004) calls a ‘double-prevention’ case: Sniper 2’s firing prevents the collision with the drone from preventing Sniper 1’s bullet striking Victim. The twist is that Sniper 2’s firing also causally suffices for Victim’s death more directly.

²³ Nor is one sufficient cause requiring the other necessary for avoiding overdetermination: witness cases of quantitative redundancy, discussed in §4.

²⁴ Gibb (2015) argues that mental events could be causally relevant to some physical event via double-prevention, without leading to systematic overdetermination. However, the kind of structure she has in mind is not like NEAR MISS; crucially, the mental event is not also causally sufficient for the physical event via an independent route, so it is more like a version of the case in which Sniper 2’s bullet sails past Victim.

²⁵ Bennett (2008: n.13) briefly mentions this strategy as a response to a similar counterexample due to Thomson-Jones (2007). An alternative strategy is pursued by Weslake (forthcoming), who adapts Bennett’s criterion using the interventionist framework. I cannot present the details here, but Weslake’s revised criterion also fails to classify cases like NEAR MISS as overdetermination.

overdetermination on downstream intermediary events (such as the two bullet-strikings) instead. Denying that the initiating causes overdetermine their effect seems counterintuitive, but it's unclear how costly this is: after all, what matters is the overall theory, and by positing downstream overdetermination we can accommodate the intuition that the overall theory is implausible (and even that the initiating causes themselves are tainted by association). This section argues against this 'intermediaries strategy'.

First, the strategy cannot be applied to cases involving direct causal processes, since there are no intermediaries on which the overdetermination can be blamed.²⁶ For example:

(DOUBLE SPELL) Working independently of each other, Morgana and Mordred each cast a spell which turns the prince into a frog at midnight. Had Mordred not cast his spell, Merlin would have cast a counter-spell to prevent the transformation.

Morgana's spell and Mordred's spell are each sufficient causes of the prince's transformation. They are on a par: the magical law gives neither precedence over the other. But Morgana's spell requires Mordred's: had Morgana cast her spell without Mordred's, Merlin would have prevented the transformation. The case mirrors NEAR MISS, with Morgana's spell mirroring Sniper 1's firing, Mordred's spell mirroring Sniper 2's firing, and Merlin's potential counter-spell mirroring the potential collision with the drone. The common structure is that some background condition needed by one sufficient cause would have failed to hold in the absence of the other. The important difference is that there are no intermediaries in DOUBLE SPELL, so the intermediaries strategy cannot be applied: *Requiring* makes the counterintuitive prediction that the prince's transformation is not overdetermined at all.²⁷

²⁶ This point parallels Schaffer's (2000: 176–7) use of 'trumping preemption' cases to argue against intermediary-based refinements of counterfactual accounts. Won (2014: 219) describes a similar counterexample.

²⁷ Variants of DOUBLE SPELL can be used as counterexamples to the other OACs above: we can parallel STARTLED SNIPER/RED FLAG with cases in which one spell-casting causes/grounds another.

Some may suspect that we cannot trust our intuitions about this kind of far-fetched case. Perhaps, for example, causation is essentially tied to locality, so that in cases involving action-at-a-distance there is no genuine causation but only some quasi-causal impostor. Even if this is right, however, I think that such cases still have dialectical force in this context. Call the relation between Morgana and Mordred's spells and the prince's transformation 'causation*'. It seems that there will be a parallel question — faced by wizard-detectives investigating the prince's transformation — of whether cases like DOUBLE SPELL exemplify an implausible kind of overdetermination*, involving distinct sufficient causes*. And it seems intuitively clear that they do: despite Morgana's spell requiring Mordred's for the prince's transformation, wizard-detectives should find the kind of causal* structure exhibited by DOUBLE SPELL *prima facie* implausible. To my mind, the failure of the requiring OAC to mitigate the implausibility of this kind of multi-determination* suggests that it also fails to mitigate the implausibility of the multi-determination in parallel real-world cases involving intermediaries, like NEAR MISS.

Examples involving direct causal processes raise an especially vivid problem for the intermediaries strategy. But my argument against it does not rely on such examples: the strategy also fails to handle other kinds of case. Consider a world containing a light that changes colour every second, with the following fundamental laws:²⁸

The Even Scarlet Law: If the light is scarlet after an even number of seconds, then it will be green the next second.

The Prime Red Law: If the light is red after a prime number of seconds, then it will be green the next second.

Now take the following case:

²⁸ By 'fundamental laws', I mean laws which do not derive their lawfulness from any other laws. (There may be fundamental laws in this sense even if, as Humeans allege, all laws ultimately derive from non-nomic facts.)

(RED LIGHT) After two seconds, the light is scarlet. After three seconds, the light turns green.

The light turning green has two sufficient causes: the light being scarlet after two seconds, in accordance with The Even Scarlet Law, and the light being red after two seconds, in accordance with The Prime Red Law. Intuitively, despite the light's being scarlet grounding the light's being red, they overdetermine the light turning green: since the laws involved are fundamental, one cause yields the effect independently of the other. In this sense, the light's turning green is caused via independent routes which happen to converge.

Through the lens of the grounding and causal relations between events, the structure of RED LIGHT mirrors that of SOPHIE (the pecking pigeon), with each featuring a 'scarlet' process and a concurrent 'red' process. In particular, in both cases — and unlike in RED FLAG — every pair of simultaneous sufficient causes is such that one grounds the other (at least, if we make the relevant assumptions about the way Sophie's visual system processes color information). Hence, proponents of *Grounding* seem forced to accept that the outputted behavior, like Sophie's pecking, is merely multi-determined. Unlike in RED FLAG, there is no pair of intermediaries on which the overdetermination can plausibly be blamed.

I conclude that the intermediaries strategy fails. Pending some other way of rescuing *Causing*, *Grounding*, and *Requiring*, compatibilists need a new way of avoiding overdetermination.

3. Diagnosis

The counterexamples above suggest a new approach for compatibilists: what matters for overdetermination is the hitherto neglected relation between the ways in which the causes are connected to their shared effect. The difference between the cases of benign multi-determination and the cases of overdetermination above seems to be that the latter involve independent 'causal connections' to the effect. The rest of the paper develops this diagnosis.

3.1 Piggybacking

In each counterexample, what seems to underlie the intuitive judgment of overdetermination is that the initiating causes are independently connected to the effect. In each of the double-sniper cases, there is the connection running via Sniper 1's bullet to Victim's death and, independently, the connection running via Sniper 2's bullet to Victim's death. In DOUBLE SPELL, there is the connection between Morgana's spell and the prince's transformation and, independently, between Mordred's spell and the prince's transformation. And in RED LIGHT, there is the connection via the Even Scarlet Law and, independently, via the Prime Red Law. In these cases, the way one cause is connected to the effect does not seem to depend on the way the other cause is connected to the effect.

In the benign multi-determination cases, by contrast, one of the causes seems to 'piggyback' on the other, in the sense that its connection to the effect depends on the other cause's connection to the effect. In BILLIARD BALLS, the connection between White's hitting Red and Black's motion is not independent of that between Red's hitting Black and Black's motion. In SOPHIE, the connection between the stimulus's being red and Sophie's pecking is not independent of that between the stimulus's being scarlet and Sophie's pecking. And in PHOTORECEPTOR, the connection between the wavelength and the current is not independent of that between the frequency and the current.

It is tempting to understand the dependence in these cases counterfactually: for example, if the connection between Red's hitting Black and Black's motion were broken, then the connection between White's hitting Red and Black's motion would also be broken. However, counterfactual dependence seems too weak to guarantee that overdetermination is avoided. Perhaps the nearest world where the causal connection between Sniper 1's firing and Victim's death is broken is one where Bodyguard jumps into the bullet's path, breaking the causal connection from Sniper 2's firing also. Even so, the two firings overdetermine Victim's death.

Instead, I propose to understand the dependence at work as a form of *metaphysical* dependence, which supports counterfactuals but may not be analyzable in terms of them. In the cases of benign overdetermination above, the connection between one sufficient cause and the effect partially

metaphysically depends on (or obtains partly in virtue of) the connection between the other sufficient cause and the effect.²⁹ This suggests the following:

Two causal connections to E do not overdetermine E if one connection metaphysically depends on the other.

This condition does not pertain to the causes themselves, but it can be converted into an OAC of the original form by the following principle: if sufficient causes C1 and C2 overdetermine E, then there are two causal connections, one from C1 to E and one from C2 to E, which overdetermine E. Say that C1 ‘piggybacks’ on C2 for E if, for every causal connection between C1 and E, CC1, and every causal connection between C2 and E, CC2, CC1 metaphysically depends on CC2. Then we have:

Piggybacking: C1 and C2 do not overdetermine E if C1 piggybacks on C2 for E.³⁰

More than one causal connection may hold between cause and effect. For example, in STARTLED SNIPER, there are two connections between Sniper 1’s firing and Victim’s death: one which proceeds via Sniper 1’s bullet, and the other which proceeds via Sniper 2’s bullet. Intuitively, the second causal connection depends on that between Sniper 2’s firing and Victim’s death, but the first does not. Hence, Sniper 1’s firing does not piggyback on Sniper 2’s firing in this case, and so *Piggybacking* does not predict that overdetermination is avoided in this case.

²⁹ In the spirit of Rosen 2010 and Fine 2012, I will not attempt to define the notion of dependence in question. It is natural to identify it with Fine’s ‘partial grounding’.

³⁰ Strictly speaking, the dependence may run in different directions for different causal connections. Say that C1 and C2 ‘co-connect’ to E if, for every causal connection between C1 and E and every causal connection between C2 and E, one connection metaphysically depends on the other. Then the principle in the text motivates the following OAC:

Co-connecting: C1 and C2 do not overdetermine E if C1 and C2 co-connect to E.

Co-connecting is stronger than *Piggybacking* (if C1 piggybacks on C2 for E, then they co-connect to E). I focus on *Piggybacking* since cases of co-connecting without piggybacking seem atypical.

In fact, once we are thinking of overdetermination in terms of causal connections, we may well deny even that it must involve distinct sufficient causes. It is natural to extend the notion so that Victim's death is overdetermined — via two independent causal connections — by Sniper 1's firing on its own, in STARTLED SNIPER, and by the scarlet flag-waving on its own, in RED FLAG. (So much for avoiding determination with a sufficiently intimate identity-like relation between causes!)

Before developing this proposal further, let me give it some general motivation. Two factors seem to underlie the *prima facie* implausibility of overdetermination. There is redundancy: we only need one of the causes to explain the effect. (We need some special reason — some evidence beyond the victim's death itself — to take Watson's double-sniper story seriously.) And there is coincidence: multiple causes converge on a single effect without good explanation. (Watson's story would be easier to believe if the two snipers were deliberately coordinating on a shared mission, rather than independently striking at the same moment.) In light of cases like STARTLED SNIPER and RED FLAG, redundancy should be understood in terms of the positing of more causal connections (rather than more sufficient causes) than is necessary to explain the effect, and coincidence should be understood in terms of the positing of causal connections (rather than sufficient causes) which converge on the effect without good explanation. Both factors are mitigated when one of the causal connections in question metaphysically depends on the other: positing them both does not count against parsimony in the same way, and there is a good explanation for their convergence. When one cause piggybacks on another, there are no independent connections, and hence no redundancy or coincidence.

The next section fleshes this OAC out by proposing an account of the central notion of 'causal connections', and the following sections show how the resulting proposal plausibly captures the cases discussed above. Although I find this particular proposal attractive, my main interest is in the general idea that *Causing*, *Grounding*, and *Requiring* should be replaced by *Piggybacking*, and more generally still, that focusing on the relation between causal connections is an important advance for compatibilists. Hence, I regard the account which follows as 'proof of concept': it serves to make the general idea both vivid and plausible. Readers who prefer an alternative account of causal connections are invited to plug it into *Piggybacking*. Given the intuitiveness of the idea

that cases of overdetermination involve independent causal connections, I view it as a desideratum for an account of causal connections that it yields the required verdicts of independence.

3.2 Causal connections

The cases in §2.4 show that the connections in question cannot be understood as mere paths of causal sufficiency (or ‘processes’), where these are sequences or mereological fusions of sufficient causes. Understanding them this way would recapitulate the mistake of the intermediaries strategy: just as the overdetermination cannot be blamed on the relations between intermediaries, so it cannot be blamed on the relations between paths constructed from these intermediaries. In RED LIGHT, for example, the ‘scarlet’ path (consisting of the light’s being scarlet throughout the second prior to its turning green) grounds the ‘red’ path (consisting of the light’s being red throughout the same second). Nonetheless, intuitively, this case involves independent connections.

Nor is the causal connection between C and E simply the fact that C is a sufficient cause for E. As I already noted, it is important that there can be more than one causal connection between C and E. In STARTLED SNIPER, [Sniper 1’s firing is a sufficient cause of Victim’s death] is plausibly partially grounded in [Sniper 2’s firing is a sufficient cause of Victim’s death], since one way in which Sniper 1’s firing suffices for Victim’s death is via Sniper 2’s firing. Nonetheless, the two firings overdetermine Victim’s death since there is another way in which Sniper 1’s firing suffices for Victim’s death which is not via Sniper 2’s firing.³¹

³¹ This problem could be avoided by defining up a more complicated sense in which one causal sufficiency fact fails to depend on the other in overdetermination cases. In particular, one might look to a notion along the following lines:

X is *exclusively grounded* via Y iff:

- i) X is partially grounded in Y
- ii) Any immediate full ground of X contains some Z which is either identical to or partially grounded in Y.

The idea would then be that, in cases like STARTLED SNIPER, one causal sufficiency fact is partially grounded in the other but is not exclusively grounded via the other (since, intuitively, it is also grounded via some separate grounding chain). More generally, piggybacking could be defined in terms of exclusive grounding as follows: C1 piggybacks on C2 for E iff [C1 is sufficient cause for E] is exclusively grounded via [C2 is a sufficient cause for E].

Reflection on the cases above suggests that the causal connections between C and E should be identified with ‘ways in which C causally suffices for E’. This may be fleshed out in terms of immediate grounding, as follows: a causal connection between C and E is a conjunction of some facts which together immediately fully ground [C is a sufficient cause for E].³² Since a single fact may have many immediate full grounds, this allows for many causal connections between C and E.

Different accounts of causal sufficiency yield different accounts of causal connections. For example, some accounts may invoke paths of energy transfer, or chains of counterfactual dependence. However, these resources seem unable to capture sufficiency on their own: one event may be connected to another by a path of energy transfer or a chain of counterfactual dependence without being causally sufficient for it. On the other hand, invoking these resources may inappropriately rule out some of the cases above: in *DOUBLE SPELL*, no path of energy transfer connects the spell-castings to the Prince’s transformation, and in the standard double-sniper case, no chain of counterfactual dependence connects the sniper’s firings to Victim’s death.

To my mind, the only plausible general approach to causal sufficiency appeals to *causal laws*: causal sufficiency is sufficiency in the circumstances given the causal laws.³³ On a schematic version of this approach, causal connections can be thought of as conjunctions consisting of three components: (i) the fact that C and E each occur, (ii) some particular background conditions B, and (iii) some general causal laws which necessitate that if C occurs under conditions B, then E occurs.³⁴ The idea is that these three components must ‘fit together’, so that they jointly necessitate

I have no particular objection to this way of developing my proposal. However, I find it more illuminating to think in terms of the immediate grounds for causal sufficiency facts, rather than these facts themselves, since this enables us to capture the idea that a single cause may have independent connections to the effect.

³² The notion of ‘immediate full ground’ is developed by Fine (2012). The existence of Socrates immediately grounds the existence of {Socrates}, and only mediately grounds the existence of {{Socrates}}.

³³ This approach is developed by Kim 1973, Mackie 1974, and, most fully, by Strevens 2007, 2008.

³⁴ I suspect that the relevant notion of ‘causal law’ is non-factive: for example, it may be a causal law that windows shatter when rocks hit them (in such-and-such circumstances) even if there are some (atypical)

that C is a sufficient cause of E. Also, they must be suitably ‘minimal’: since the components ground causal sufficiency facts, causal connections must not contain any redundant or irrelevant causal laws or background conditions. But if they do fit together in a suitably minimal way, then the three components jointly constitute what I am calling a ‘causal connection’ between C and E: they jointly give rise to the fact that C is causally sufficient for E.

Of course, this is very far from being a complete account as it stands: there is a lot of philosophical work to be done in accounting for the crucial notions of ‘causal law’ and ‘background condition’ (and the corresponding notion of ‘minimality’). But for my purposes, this schematic version will be enough to get on with: we have a sufficient grasp of relevant causal laws and background conditions to make the needed judgments of dependence between causal connections so understood.

In BILLIARD BALLS, for example, the causal sufficiency of Red’s hitting Black for Black’s motion holds in virtue of some general laws of dynamics (such as the conservation of momentum) together with some particular facts concerning Red and Black’s situation (such as the fact that Black is not stuck to the table). Call the conjunction of these general laws and particular facts, together with the fact that Red hit Black and Black moved, ‘*Red-Black*’; it is a causal connection between Red’s hitting Black and Black’s motion. Similarly, the causal sufficiency of White’s hitting Red for Red’s hitting Black holds in virtue of some general laws of dynamics (perhaps the same as in *Red-Black*) together with some particular facts concerning White, Red and Black’s situation (such as Red’s initial position relative to Black). Call the conjunction of these general laws and particular facts, together with the fact that White hit Red and Red hit Black, ‘*White-Red*’; it is a causal connection between White’s hitting Red and Red’s hitting Black.

exceptions. If so, then what necessitates is the law’s content, rather than the fact that it is a causal law. There are important and difficult questions in the vicinity. For example, how should the account be generalized to cases of probabilistic causation? Any attempt to properly address such questions would take us too far from the main thread. The point is not to defend a fleshed-out account of causal sufficiency, but to argue that, on a plausible general approach to causal sufficiency, the piggybacking OAC is vindicated.

3.3 Avoiding overdetermination

In this section, I motivate *Piggybacking* — with causal connections understood in the way I have outlined — by showing how it captures what is benign about the cases of benign multi-determination above. In the next section, I show how it avoids the counterexamples to the other OACs.

Two useful though defeasible heuristics support my verdicts about metaphysical dependence in this section and the next. First, the manipulability test: if A metaphysically depends on B, then (ceteris paribus) preventing B from obtaining is a way to prevent A from obtaining. For example, if the temperature increase in the room metaphysically depends on the increasing kinetic energy of the air molecules in the room, then preventing their kinetic energy from increasing is a way to prevent the temperature from increasing. Conversely, since preventing the kinetic energy of some faraway air molecules from increasing does not prevent the temperature increase inside the room, the temperature increase does not metaphysically depend on the increasing kinetic energy of these faraway molecules.³⁵ Second, the explanation test: if A metaphysically depends on B, then (ceteris paribus) B's obtaining partially metaphysically (as opposed to causally) explains A's obtaining. For example, the increasing kinetic energy of the air molecules in the room partially metaphysically explains the temperature increase in the room, whereas the increasing kinetic energy of the faraway air molecules does not.³⁶

With these tests in mind, I will argue that there is metaphysical dependence between the causal connections in the cases of benign multi-determination above, and that there is not in the cases of overdetermination.


³⁵ Ceteris are not paribus when B is prevented in a way which either replaces it with or leaves in place something else which determines A. In the cases of independent connections discussed below, the failure to pass the manipulability test is not plausibly attributable to reasons like these.

³⁶ Ceteris are not paribus when (for pragmatic reasons) connecting A to B fails to alleviate puzzlement (e.g. because the explanation's audience is incapable of understanding the connection, or because their puzzlement pertains to B as much as to A). In the cases of independent connections discussed below, the failure to pass the explanation test is not plausibly attributable to reasons like these.

In BILLIARD BALLS, the causal connection from Red’s hitting Black to Black’s motion is *Red-Black*: the conjunction of those general causal laws and particular background conditions which immediately ground the causal sufficiency of Red’s hitting Black for Black’s motion, together with the fact that Red’s hitting Black and Black’s motion each occurred. The causal connection from White’s hitting Red to Black’s motion is the conjunction of those general causal laws and particular background conditions which immediately ground the causal sufficiency of White’s hitting Red for Black’s motion, together with the fact that White’s hitting Red and Black’s motion each occurred. Call this conjunction ‘*White-Black*’.

It is plausible that *White-Black* metaphysically depends on *Red-Black*: intuitively, White’s hitting Red is connected to Black’s motion *via* the connection between Red’s hitting Black and Black’s motion. Applying the manipulability test: preventing *Red-Black* from obtaining (e.g. by preventing Red from hitting Black, or by supergluing Black to the table), is a way of preventing *White-Black* from obtaining. Applying the explanation test: *White-Black* obtains partly because *Red-Black* obtains (where this partial explanation may be completed by adding that *White-Red* — the connection between White’s hitting Red and Red’s hitting Black — also obtains).

The situation is illustrated below, with single-tailed arrows representing causal connections and double-tailed arrows representing metaphysical dependence:



More generally, it is plausible that any causal connection between a mediate (or ‘distal’) cause and its effect holds partly in virtue of the sub-connections between its intermediaries and the effect. (Of course, as cases like STARTLED SNIPER illustrate, this is consistent with a mediate cause and one of its intermediaries also being independently connected to their shared effect.)


In SOPHIE, the causal connection from the red stimulus's presentation to Sophie's pecking is the conjunction of the causal laws and background conditions which ground the causal sufficiency of the red stimulus's presentation for Sophie's pecking, together with the fact that each of these events occurred. The relevant general laws relate Sophie's being presented with a red stimulus to her pecking, and the relevant background conditions presumably include that Sophie was not blindfolded, that she was not prevented from pecking in any way, and so on. Call this conjunction '*Red-Pecking*'.

The causal connection from the scarlet stimulus's presentation to Sophie's pecking is the conjunction of the causal laws and background conditions which ground the causal sufficiency of the scarlet stimulus's presentation for Sophie's pecking, together with the fact that each of these events occurred. The relevant general laws relate Sophie's being presented with a scarlet stimulus to her pecking, and the relevant background conditions are presumably the same as those in *Red-Pecking* (perhaps without any contingencies which are only needed to ensure that Sophie pecks in situations where the stimulus is red but not scarlet). Call this conjunction '*Scarlet-Pecking*'.³⁷

It is plausible that *Red-Pecking* metaphysically depends on *Scarlet-Pecking*: intuitively, the overall connection between red and pecking depends on the specific connection between scarlet and pecking. Applying the manipulability test: preventing *Scarlet-Pecking* from obtaining (e.g. by preventing the presentation of the scarlet stimulus, or by rendering Sophie incapable of recognizing scarlet), is a way of preventing *Red-Pecking* from obtaining. Applying the explanation test: *Red-Pecking* obtains partly because *Scarlet-Pecking* obtains (where this partial explanation is completed by adding any laws and background conditions need to ensure that non-scarlet shades of red would also have sufficed for pecking).

The situation is illustrated below:

³⁷ Perhaps the very same laws and background conditions which ground the causal sufficiency of red also ground that of scarlet. Then the dependence of *Red-Pecking* on *Scarlet-Pecking* comes down to the dependence of the red stimulus's presentation on the scarlet stimulus's presentation.



More generally, it is plausible that the causal connection between a determinable cause and its effect holds partly in virtue of the causal connections between its specific determinates and the effect. (Of course, as cases like RED FLAG and RED LIGHT illustrate, this is consistent with a determinable cause and one of its determinates also being independently connected to their shared effect.)³⁸

In PHOTORECEPTOR, the causal connection from the frequency of the incoming light to the produced current is the conjunction of the causal laws and background circumstances which ground the causal sufficiency of the frequency for the current, together with the fact that these two events occurred. The relevant general laws relate the incoming frequency to the current produced, and the relevant background conditions include that the photoreceptor is wired up correctly. Call this conjunction *Frequency-Current*.


The causal connection from the wavelength of the incoming light to the produced current is the conjunction of the causal laws and background circumstances which ground the causal sufficiency of the wavelength for the current, together with the fact that these two events occurred. The relevant general laws relate the incoming wavelength to the current produced, with the relevant background conditions including that the photoreceptor is wired up correctly and that it is

³⁸ Lee (2021) argues that the fact that the red stimulus causes pecking partially grounds the fact that the scarlet stimulus causes pecking since, intuitively, the latter causal fact derives from chaining together the fact that the scarlet stimulus grounds the red stimulus with the fact that the red stimulus causes pecking. Although Lee focuses on the causal facts themselves rather than causal connections, an analogous view may be tempting: perhaps *Scarlet-Pecking* depends on *Red-Pecking* together with the ‘metaphysical connection’ between the scarlet stimulus and the red stimulus. (This would make the case similar to that of BILLIARD BALLS, with one of the causal connections replaced by a metaphysical connection.) Though I prefer the direction of dependence presented in the text, reversing the dependence would still vindicate the main claim that one cause piggybacks on the other in cases like SOPHIE.

surrounded by air (in a different medium, wavelengths correspond to different frequencies). Call this conjunction *Wavelength-Current*.

It is plausible that *Wavelength-Current* metaphysically depends on *Frequency-Current*: intuitively, the wavelength connects to the current *via* the connection between the frequency and the current. Applying the manipulability test: preventing *Frequency-Current* from obtaining (e.g. by changing the frequency of the incoming light, or by interfering with the photoreceptor’s wiring), is a way of preventing *Wavelength-Current* from obtaining. Applying the explanation test: *Wavelength-Current* obtains partly because *Frequency-Current* obtains (where this partial explanation is completed by adding the synchronic nomic connection between the wavelength and the frequency).

The situation is illustrated below, with the dashed arrow representing a synchronic nomic connection:



More generally, whenever two events are nomically connected, one of them may inherit a causal connection to some third event from the other’s causal connection to that third event: piggybacking may be mediated by a synchronic nomic connection.³⁹

³⁹ Naturalistic dualists might view mental causation as a case of this kind. In particular, they may hold that there is a causal connection between mental state M and some behavior B in virtue of i) M’s nomic connection to some physical state P, and ii) P’s causal connection to B. If so, the causal connection between M and B metaphysically depends on that between P and B, and hence overdetermination is avoided by M piggybacking on P for B.

Some may find it objectionable that *Piggybacking* accommodates dualist mental causation without overdetermination. But regarding this case as a datum puts the cart before the horse: we need an OAC which works in pre-theoretically intuitive cases before we can apply it to dualism. It is worth noting also that even

3.4 Avoiding the counterexamples

So far, I have shown that *Piggybacking* provides a unified explanation of why the multi-determination in BILLIARD BALLS, SOPHIE, and PHOTORECEPTOR is benign: in each case, one causal connection metaphysically depends on the other. It remains to show that *Piggybacking* does not fall victim to the counterexamples which jeopardize *Causing*, *Grounding*, and *Requiring*.

By contrast to the cases of benign multi-determination, analogous metaphysical dependence claims are implausible in the cases of overdetermination presented above. In STARTLED SNIPER and NEAR MISS, there is a causal connection between Sniper 1's firing and Victim's death, *Sniper 1-Victim*, which is independent of that between Sniper 2's firing and Victim's death, *Sniper 2-Victim*. Perhaps both causal connections involve the same general laws: a ballistics law connecting sniper-firings to bullet-strikings, and a physiological law connecting bullet-strikings to deaths. They also share some particular background conditions, such as the fact that Victim is not wearing a bullet-proof vest. But some of the background conditions involved, and of course the causes themselves, differ: *Sniper 1-Victim* includes that Sniper 1 is within range, with a clear line of sight, whereas *Sniper 2-Victim* includes the corresponding facts about Sniper 2.

Sniper 1-Victim does not plausibly depend on *Sniper 2-Victim*, or vice versa; intuitively, neither connection flows through the other. Applying the manipulability test: there are ways of preventing *Sniper 2-Victim* from obtaining (e.g. by placing an obstacle between Sniper 2 and Victim), which are not ways of preventing *Sniper 1-Victim* from obtaining. Applying the explanation test: *Sniper 2-Victim* does not even partly explain *Sniper 1-Victim*. In particular, Sniper 2's firing, and the background conditions specific to *Sniper 2-Victim*, seem explanatorily irrelevant to *Sniper 1-Victim*.⁴⁰

if dualism has no problem with overdetermination *per se*, this doesn't imply that it has no problem in the vicinity (for example, perhaps it involves objectionable redundancy in the fundamental laws).

⁴⁰ At least, in the metaphysical sense of explanation that is relevant to the metaphysical dependence test. Causally speaking, Sniper 2's firing does help to explain Sniper 1's firing in STARTLED SNIPER, and one of the background conditions in NEAR MISS (the absence of a colliding drone).

In RED FLAG, there is a causal connection between the red flag-waving and Victim's death, *Red-Victim*, going via Sniper 2's firing, and a causal connection between the scarlet flag-waving and Victim's death, *Scarlet-Victim*, going via Sniper 1's firing. In addition to the laws and background conditions involved in *Sniper 1-Victim/Sniper 2-Victim*, these causal connections involve further laws and background conditions which ensure that the waving of a scarlet/red flag is sufficient for Sniper 1/Sniper 2 firing respectively.

Red-Victim does not plausibly depend on *Scarlet-Victim* (or vice versa); intuitively, neither connection flows through the other. Applying the manipulability test: there are ways of preventing *Scarlet-Victim* from obtaining (e.g. by preventing Sniper 1 from seeing the flag-waving), which are not ways of preventing *Red-Victim* from obtaining.⁴¹ Applying the explanation test: *Red-Victim* does not even partly explain *Scarlet-Victim*. In particular, the general laws and background conditions specific to *Scarlet-Victim* (which connect the scarlet flag-waving to Sniper 1's firing), seem explanatorily irrelevant to *Red-Victim*.

Similarly, in DOUBLE SPELL, there is a causal connection between Morgana's spell and the prince's transformation, *Morgana-Prince*, which is independent of that between Mordred's spell and the prince's transformation, *Mordred-Prince*. These causal connections presumably involve the same general law: the law of magic corresponding to the transformation spell that Morgana and Mordred each use. But the causes themselves — and perhaps also the relevant background conditions — differ.

Morgana-Prince does not plausibly depend on *Mordred-Prince*, or vice versa; intuitively, neither connection flows through the other. Applying the manipulability test: there are ways of preventing *Mordred-Prince* from obtaining (e.g. by casting a spell which renders Mordred's spells inefficacious) which are not ways of preventing *Morgana-Prince* from obtaining. Applying the explanation test: *Mordred-Prince* does not even partly explain *Morgana-Prince* (or vice versa). In

⁴¹ Another way to disrupt *Scarlet-Victim* without disrupting *Red-Victim* is to make the flag a different shade of red. However, this fails to show that the latter does not depend on the former, since the test fails analogously in SOPHIE. This is one of those cases (mentioned in n.35) where *ceteris are not paribus*. (The same point applies to RED LIGHT.)

particular, Mordred's spell-casting (and any background conditions specific to *Mordred-Prince*, such as Mordred's not being subject to any impotency spell) seem explanatorily irrelevant to *Morgana-Prince*.⁴²

Finally, in RED LIGHT, there is a causal connection between the light's being scarlet and its turning green, *Scarlet-Green*, and an independent causal connection between the light's being red and its turning green, *Red-Green*. In this case, the fundamental laws directly connect the colour of the light at different times, so that no background conditions are involved. But the fundamental laws corresponding to the two causal connections (as well as the causes themselves) are different. In particular, *Scarlet-Green* involves the Even Scarlet Law, whereas *Red-Green* involves the Prime Red Law.

Red-Green does not plausibly depend on *Scarlet-Green*, or vice versa; intuitively, neither connection flows through the other. Applying the manipulability test: there are ways of preventing *Scarlet-Green* from obtaining (namely, by deleting or modifying the Even Scarlet Law) which are not ways of preventing *Red-Green* (since they leave the Prime Red Law intact). Applying the explanation test: *Scarlet-Green* does not even partly explain *Red-Green* (or vice versa). Since each law is fundamental, neither is explanatorily relevant to the other.

I conclude that *Piggybacking* captures what is benign about all the cases of benign multi-determination above, while avoiding the counterexamples that undermined the OACs in §2. Compatibilists can make progress by focusing on causal connections.

⁴² At least, in the metaphysical sense of explanation that is relevant to the metaphysical dependence test. Causally speaking, Mordred's spell-casting does help to explain one of the background conditions in *Morgana-Prince*: namely, the fact that Merlin did not cast a counter-spell.

4. Conclusion

Compatibilist responses to causal exclusion arguments turn on the conditions under which overdetermination is avoided. One such condition is that one cause piggybacks on the other. This unifies a broad class of cases of benign multi-determination, including all those presented in §2.

It does not cover all cases, however. Consider *quantitative redundancy*: a rock shattering a window is decomposable into two independently sufficient half-rocks (Schaffer 2003: 28), a sledgehammer flattening a chestnut is decomposable into two independently sufficient half-sledgehammers (Mackie 1974: 43), and so on. Such cases seem to involve independent causal connections, but they are not overdetermination (in my sense): multi-determination of this kind is inevitable, and commitment to it does not jeopardize a theory. Thus, piggybacking is not necessary for avoiding overdetermination.

This raises the question of what other OACs there may be, and whether an exhaustive list would provide a unified account of what overdetermination *is*. Whilst I cannot properly address this question here, let me close by outlining my own approach to it.

In cases of quantitative redundancy, the converging causal connections seem to be explainable as abstractions from the ‘whole’ causal connection which encompasses each of them. For example, the connection between the left half-rock hitting the window and the window shattering and the connection between the right half-rock hitting the window and the window shattering each seem explainable as ‘abstractions’ from the connection between the rock hitting the window and the window shattering. This suggests that even if the converging connections do not depend on each

other, they may nonetheless fail to be independent in the sense that they each depend on some common connection.⁴³

In many cases, a mechanism has been selected to produce an effect via independently operating (and hence, usefully redundant) causal connections. Each engine's functioning is causally sufficient for the plane's flight, each kidney's functioning is causally sufficient for our blood being filtered, and so on. This supports a teleological explanation for the redundancy: the existence of multiple independent causal connections is explained by means of its function (via some form of selection process).

What these cases seem to share with the cases of benign multi-determination in §2 is that the convergence of distinct causal connections can be satisfactorily explained, without jeopardizing the overall theory. Sometimes this is because one connection metaphysically depends on — and hence can be explained in terms of — the other, but many other forms of explanatory mitigation are conceivable.

If this approach is correct, asking what it is for an effect to be overdetermined is like asking what it is for a posit to be 'ad hoc'. Overdetermination has its proper home in the epistemology of theory choice rather than the metaphysics of causation: our grip on it is via its role in making theories epistemically implausible rather than metaphysically impossible. From this perspective, exclusion arguments are a particular application of inference to the best explanation. Avoiding overdetermination is ultimately about providing a good explanation, where this is a (somewhat indeterminate and context-dependent) matter of a theory's overall explanatory virtues given our evidence. The only general theory of overdetermination to be had is that multi-determination is overdetermination when the converging causal connections tarnish a theory's explanatory goodness.

⁴³ This is similar to Paul's (2007) idea that an object and its constituting matter can each be causes without overdetermination since they each cause in virtue of a single shared property-instance.

Acknowledgements

For discussion and comments at various stages, thanks to Karen Bennett, Verónica Gómez, Barry Loewer, David Papineau, Andrew Rubner, Jonathan Schaffer, Ted Sider, Bram Vaassen, Isaac Wilhelm, the Spring 2020 Rutgers Dissertation Seminar, two anonymous referees for this journal, and several anonymous referees for other journals.

References

- Aimar, S. (2011). Counterfactuals, Overdetermination and Mental Causation. *Proceedings of the Aristotelian Society* (Hardback), 111: 469–477. doi:10.1111/j.1467-9264.2011.00318.x
- Albert, D. (2015). *After Physics*. Harvard University Press.
- Bennett, K. (2003). Why the Exclusion Problem Seems Intractable, and How, Just Maybe, to Tract It. *Noûs*, 37: 471–97.
- (2008). Exclusion again. In J. Hohwy & J. Kallestrup (eds.), *Being Reduced: New Essays on Reduction, Explanation, and Causation*. Oxford University Press. pp. 280–307.
- Bernstein, S. (2016). Overdetermination Underdetermined. *Erkenntnis* 81 (1):17-40.
- Chalmers, D. (2013). Panpsychism and Panprotopsychism. *Amherst Lecture in Philosophy* 8.
- Crane, T. & Árnadóttir, S. T. (2013). There is No Exclusion Problem. In E. J. Lowe, S. Gibb & R. D. Ingthorsson (eds.), *Mental Causation and Ontology*. Oxford: Oxford University Press, pp. 248-66.
- 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: Oxford University Press.
- Fine, K. (2012). Guide to Ground. In F. Correia & B. Schnieder (eds.), *Metaphysical Grounding*. Cambridge University Press, pp.37–80.
- Gibb, S. (2015). VIII—Defending Dualism. *Proceedings of the Aristotelian Society* 115 (2pt2):131–146.

- Goldman, A. (1969). The Compatibility of Mechanism and Purpose. *The Philosophical Review* 78: 468–482.
- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall & L.A. Paul (eds.), *Causation and Counterfactuals*. MIT Press. pp. 225–276.
- Kim, J. (1973). Causation, nomic subsumption, and the concept of event. *Journal of Philosophy* 70 (8): 217–236.
- (1993). The Nonreductivist’s Troubles with Mental Causation. Reprinted (1993) in *Supervenience and Mind*. Cambridge: Cambridge University Press, 336–357.
- (1998). *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge, MA: MIT Press.
- Kroedel, T. & Schulz, M. (2016). Grounding mental causation. *Synthese* 193(6): 1909–1923.
- Lee, S. (2021). Building low level causation out of high level causation. *Synthese* 199(3-4):9927–9955.
- Lepore, E. & Loewer, B. (1987). Mind Matters. *The Journal of Philosophy* 84: 630–642.
- Lewis, D. (1973). Causation. Reprinted (1986) in *Philosophical Papers, Volume II*. Oxford: Oxford University Press. 159–213.
- List, C. & Menzies, P. (2009). Nonreductive Physicalism and the Limits of the Exclusion Principle. *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.
- Mackie, J. L. (1974). *The Cement of the Universe*. Oxford: Oxford University Press.
- Mills, E. (1996). Interactionism and Overdetermination. *American Philosophical Quarterly* 33: 105–117.
- Papineau, D. (2001). The rise of physicalism. In C. Gillett & B. Loewer (eds.), *Physicalism and its Discontents*. Cambridge: Cambridge University Press.
- (2002). *Thinking About Consciousness*. Oxford: Oxford University Press.

- Paul, L. A. (2007). Constitutive Overdetermination. In J. K. Campbell, M. O'Rourke & H. Silverstein (eds.), *Causation and Explanation*. Cambridge: MIT Press. pp. 265—290.
- Pereboom, D. (2002). Robust nonreductive materialism. *The Journal of Philosophy* 99 (10): 499–531.
- Rescorla, M. (2014). The Causal Relevance of Content to Computation. *Philosophy and Phenomenological Research* 88 (1): 173–208.
- Rosen, G. (2010). Metaphysical Dependence: Grounding and Reduction. In B. Hale & A. Hoffmann (eds.), *Modality: Metaphysics, Logic, and Epistemology*. Oxford: Oxford University Press.
- Schaffer, J. (2000). Trumping preemption. *The Journal of Philosophy* 97(4): 165–181.
- (2003). Overdetermining causes. *Philosophical Studies* 114: 23–45.
- (2015). What Not to Multiply Without Necessity. *Australasian Journal of Philosophy* 93(4): 644–664.
- Shapiro, L. A. (2012). Mental Manipulations and the Problem of Causal Exclusion. *Australasian Journal of Philosophy* 90 (3): 507–524.
- Shapiro, L. A. & Sober, E. (2007). Epiphenomenalism - the do's and the don'ts. In G. Wolters & P. K. Machamer (eds.), *Thinking About Causes: From Greek Philosophy to Modern physics*. University of Pittsburgh Press. pp. 235–264.
- Shoemaker, S. (2001). Realization and mental causation. In C. Gillett & B. Loewer (eds.), *Physicalism and its Discontents*. Cambridge: Cambridge University Press.
- Sider, T. (2003). What's So Bad About Overdetermination? *Philosophy and Phenomenological Research* 67 (3):719–726.
- Strevens, M. (2007). Mackie Remixed. In J. K. Campbell, M. O'Rourke & H. Silverstein (eds.), *Causation and Explanation*. Bradford. pp. 93–118.
- (2008). *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- Thomson-Jones, M. (2007). Overdetermination, Mental Causation, and the Staggered Firing Squad: Problems with the Case for Compatibilism. Unpublished manuscript.

- Vaassen, B. (2021). Causal Exclusion without Causal Sufficiency. *Synthese* 198: 10341–10353.
- Weslake, B. (forthcoming). Exclusion Excluded. In A. Wilson & K. Robertson (eds.), *Levels of Explanation*. Oxford University Press.
- Wilson, J. (2009). Determination, realization and mental causation. *Philosophical Studies* 145(1): 149–169.
- Won, C. (2014). Overdetermination, Counterfactuals, and Mental Causation. *Philosophical Review* 123 (2): 205-229.
- Woodward, J. (2008). Mental causation and neural mechanisms. In J. Hohwy & J. Kallestrup (eds.), *Being Reduced: New Essays on Reduction, Explanation, and Causation*. Oxford: Oxford University Press. pp. 218-262.
- (2015). Interventionism and Causal Exclusion. *Philosophy and Phenomenological Research* 91 (2):303-347.
- Yablo, S. (1992). Mental causation. *Philosophical Review* 101:245–80.
- Zhong, L. (2020). Taking Emergentism Seriously. *Australasian Journal of Philosophy* 98 (1):31-46.