Philosophers' Imprint

VOLUME 25, NO. 43
DECEMBER 2025

HOW MUCH DID EACH OF THE AUTHORS OF THIS PAPER CAUSALLY CONTRIBUTE TO ITS WRITING?

Sam Baron, Helen Beebee, David Braddon-Mitchell, Antony Eagle, and Kristie Miller

University of Melbourne, University of Leeds/University of Helsinki, University of Sydney, University of Adelaide, and University of Sydney

doi.org/10.3998/phimp.6191

1. Introduction

Philosophers sometimes claim that causation or 'causal contribution' is something there can be more or less of (Mumford 2013, Braham and van Hees 2009, Sprenger 2018, Demirtas 2022, Bernstein 2017, Moore 2009, Goh 2023, Mumford and Anjum 2011, Northcott 2005). They claim that it is perfectly sensible to ask whether nationalism was more of a cause of the First World War than was militarism, or whether *A* made more of a causal contribution to the writing of the paper than did *B*. Some such authors go further, claiming that either relative or absolute causal contribution is measurable: it makes sense to ask whether *A*'s causal contribution was twice that of *B*, or whether *A*'s contribution was greater or less than 23%, and so on. And, of course, we might expect perfectly sensible questions to have perfectly sensible answers. Presumably in these kinds of case the correct answer is whatever is delivered by the correct theory of degree of causation or causal contribution. We'll call such theories 'degree-theoretic' theories, and we'll call the general view that the extent or degree of causation or causal contribution is measurable a 'degree-theoretic' conception of causation.

According to Mumford "the scalarity of causes must be accommodated" (2013: 111). On the contrary, we will argue that the degree-theoretic conception of causation is misguided. None of the reasons philosophers have adduced for thinking that a degree-theoretic conception is needed are good reasons. Moreover, different degree-theoretic theories deliver very different results—and, while a given theory's results might do the job required by some particular alleged need, it is unclear why we should privilege one such alleged need over the others, and hence unclear how we might rationally decide between the rival degree-theoretic theories.

Our first claim is this: we can happily grant that the demand for a degree-theoretic conception of causation is a legitimate demand for *something* that comes in degrees, but it is not necessary that what needs to come in degrees is causation (or causal contribution). The various features of the world that different degree-theoretic theories take to

deliver degrees of causation are genuine features of the world. So, we can appeal directly to the relevant feature of the world to do whatever must be done. There is, then, no work for the notion of degree-theoretic causation to play and so positing such causation is unnecessary.

Our second claim is this: not only is it unnecessary to posit degree-theoretic causation but degree-theoretic causation cannot in fact play the roles that motivate those who posit it. Thus, insofar as one is motivated to posit degree-theoretic causation because one sees that it is important that these roles are played, one shouldn't posit degree-theoretic causation as the thing that plays those roles.

A look ahead. In §2, we consider the various reasons offered in the literature for thinking that causation comes in degrees, and we classify them into three arguments, and briefly justify our later focus on just one of them, 'the Practice Argument'. In §3, we provide a (non-exhaustive) survey of theories of degree of causation and briefly say something about how and why they deliver very different results from one another. In §4, we argue that the Practice Argument fails to establish that causation comes in degrees. In §5 we argue not only that causation *needn't* be understood to come in degrees, but that it *oughtn't*.

2. Arguments in favour of a degree-theoretic conception of causation

Before proceeding, it's worth offering two clarifications. First, our concern in this paper is with 'token' or 'actual' causation: causation between particular events (or whatever your preferred relata are), and not causation at the 'type' level. For what it's worth, the claim that type-causation comes in degrees—or at least relative strength (as in 'smoking is a stronger cause of heart disease than is obesity in the UK population')—strikes us as plausible, and nothing we say is incompatible with it.

Second, we have so far been talking about (purported) degrees of 'causation, or causal contribution', and this distinction requires explanation. A relation can be collective (as in 'Bob and Alice surrounded the tree'—something neither of them individually does) or distributive (as in 'the orange and the banana are sitting on the table'—something they each individually do). We can make the same distinction when it comes to causation. We could take the sentence 'x and y caused z' to imply that x individually caused z, and so did y: that's the distributive reading. Or (setting aside cases of overdetermination) we could take it to imply that x and y *collectively* caused z.

Many theories of causation take the distributive approach; on a distributive view, 'x caused z' and 'x was a cause of z' are synonyms. By contrast, Alex Kaiserman (2016) takes a collective approach: to say that x is *a* cause of z is to say that x is *one* of the several events that *collectively* caused z. Kaiserman claims that it makes no more sense to say that one thing was more of a cause than another, or that x caused z to a greater degree than did y, than it does to say that Alice surrounded the tree more than Bob did. In his view causation doesn't come in degrees: either x was a cause of z or it wasn't. On the other hand, it *does* make sense to say that Alice contributed more to the surrounding of the tree than Bob did—assuming they surrounded the tree collectively rather than distributively. Similarly, it makes sense to say that x *contributed more to the causing* of z than did y—assuming that x and y are amongst the collective causes of z.

We have both views in our sights, and so we intend 'degree-theoretic theories of causation' to be read as neutral with respect to whether causation is collective or distributive, and hence neutral with respect to whether it is causation (distributive) or causal contribution (collective) that is alleged to be a matter of degree.

With this in mind, why think causation is degree-theoretic? Three sorts of argument have been offered. The first is the *entailment argument*: that causation comes in degrees is simply entailed by the correct theory of causation (e.g., Lewis's (2000) theory).

The Entailment Argument

(1) According to theory *T*, causation comes in degrees.
(2) Theory *T* is the correct theory of causation.
Therefore,
(3) Causation comes in degrees.

We will have little to say in this paper about this style of argument. We are predominantly interested in arguments that try to show that we *should* endorse a theory of causation that is at the very least consistent with causation coming in degrees, and that we should do this because there is independent reason to think that causation does come in degrees.

So, let's consider arguments of this type. First, the *Intuition Argument*:

The Intuition Argument

(4) It is intuitive that causation comes in degrees.
(5) If it is intuitive that causation comes in degrees, then causation comes in degrees.
Therefore,
(6) Causation comes in degrees.

Why think (4) is true? Several philosophers have taken empirical studies of people's causal judgements as evidence that people conceive of causation as coming in degrees, and hence as reason to suppose that it does (Sprenger 2018, Goh 2023). As O'Neill *et al*. (2022) note, psychological experiments that seek to elicit people's causal judgements often ask participants to respond to a question such as "To what extent do you agree with the statement '*X* caused *Y*?' " on a Likert scale ranging from "strongly disagree" to "strongly agree" (e.g., Icard *et al*. 2017; Kominsky *et al*. 2015; Henne *et al*. 2021), while some even explicitly ask "To what degree did *X* cause/prevent *Y*?" (Gerstenberg *et al*. 2021; Lagnado and Channon 2008). If people do not take causation to be degree-theoretic we would expect them to use the two ends of the

Likert scale to report that there is, or is not, causation present. In fact, though, participants not only respond at the scale extremes but also at most values between those extremes. Indeed, in Icard *et al*. (2017) and Morris *et al*. (2019), causal judgments at the center of the scale were about as frequent as responses at the scale extremes.

There are, however, good reasons to resist (4). First, participants in experiments tend to interpret questions charitably, and do their best to provide some sort of sensible answer. It could be that participants simply read such questions as asking them not about degrees of causation, but rather about their degree of belief that the event is a cause. This has been suggested by a number of authors (see, e.g., Collins and Shanks 2006; Liljeholm and Cheng 2009; Schlottmann and Anderson 1993). Indeed, evidence suggests that people's degree-theoretic causal judgments and their confidence in those judgments align (Liljeholm and Cheng 2009), suggesting that they may be using the Likert scale to report their degrees of belief rather than their beliefs in degree-theoretic causation. Recent work by O'Neill *et al*. (2022) suggests this might be so. If the gradation in causal judgements reflects uncertainty about causation, rather than a judgement about degree of causation, we would expect to find that graded causal judgments are made with low confidence, whereas those made at the extreme ends of the scale are made with high confidence. O'Neill *et al*. indeed found precisely this expected pattern. As we see it, then, this sort of experimental work does not give us reason to think that people intuitively suppose causation to be degree-theoretic. In what follows then, we set this argument aside.

Next, is *the Practice Argument*:

The Practice Argument

(7) Practice *P* is reasonable.
(8) Practice *P* is reasonable only if causation comes in degrees.
Therefore,
(9) Causation comes in degrees.

The Practice Argument can be filled out with various different sorts of practices. Three practices seem especially relevant: moral responsibility practices, legal responsibility practices, and scientific practices.

Consider first our practices surrounding moral responsibility. Those practices involve holding people more, or less, morally culpable for certain outcomes. For instance, suppose there is a car crash in which *F* is the driver of one car, and *A* of the other. Suppose that *F* is busy texting with one hand, while *A* is paying careful attention to where she is going. We might well think that *F* is *more* culpable for the resulting crash than is *A*. Quite generally we have a host of practices surrounding moral responsibility in which we differentially praise, blame, reward and punish. These practices seem entirely reasonable. But many philosophers have thought that the only way to make sense of the reasonableness of these practices is to suppose that causation itself is degree-theoretic. For instance, Bernstein says that "having an adequate theory of such *degrees* of causal contribution is essential for unpacking the role of causation in attributions of moral responsibility" (Bernstein 2017: 1). Braham and van Hees (2009) and Moore (2009) make similar claims. The idea, we take it, is that moral responsibility tracks causal 'responsibility' and so if moral responsibility comes in degrees, then this must be because causation comes in degrees. The reason that *F* is more culpable for the crash is that he is more of a cause of the crash than *A*.

Next, consider our legal practices. Suppose that *A*'s and *F*'s cars crashed into *M*'s house. *M* has no insurance, and she sues *A* and *F* for compensation in order to rebuild her front room. Even if both *F* and *A* are legally responsible for the crash, and hence liable to compensate *M*, the law may well say that *F* is more legally responsible for the crash, and hence must compensate *M* more than *A* must. Quite generally, legal responsibility is taken to be a matter of degree. And its being so seems reasonable. But many philosophers have thought that the only way we can make sense of legal responsibility being a matter of degree is if causation is a matter of degree (Moore 2009, Demirtas 2022, Braham and van Hees 2009). Again, we take it that the idea is that legal responsibility simply tracks causal responsibility. It is because causation comes in degrees that legal responsibility does. As Moore puts it, "given the scalarity of causation, all the law need do is draw the line for liability somewhere on the scale of causal contribution" (Moore 2009: 105).

Finally, consider our scientific practices. Scientific explanations often involve reporting the relative weight of different factors in accounting for some outcome. We wonder which factors contributed more to the pollution of the river; we wonder whether it is genetic or environmental factors that contribute more to cancer; we wonder which country, politician or policy contributed the most to the COVID-19 pandemic death toll (Demirtas 2022, Mumford and Anjum 2011, Northcott 2005). These practices of finding and weighting such factors seems reasonable. But, the thought goes, in tracking the relative importance of these factors we just are tracking their causal contribution. When we ask about the relative roles of this or that policy in the spread of COVID-19, we must in fact be asking which of those policies was more of a cause of the spread of the infection (Demirtas 2022). Thus, if these practices are reasonable, it must be because causation comes in degrees.

In what follows we will grant that each of these different practices is reasonable. We will argue, however, that their reasonableness does *not* require that causation comes in degrees. So, we accept versions of (7) and deny versions of (8). We make this case in §4. First, however, we will briefly survey a number of degree-theoretic accounts of causation. This is important, since those accounts draw on other degree-theoretic phenomena, and our aim in §4 will be to argue that such degree-theoretic phenomena can be put to work in vindicating the reasonableness of our practices—it is just not causation.

3. A survey of theories of degree-theoretic causation

There are several proposals available for how we specify degrees of causation. Without trying to be exhaustive, we will briefly outline some of the main options (for more detailed surveys, see Demirtas 2020 and

Sartorio 2022). Two main kinds of approach take degree of causation to be a matter of either degree of necessity (see e.g. Northcott 2005, 2008a, 2008b) or degree of sufficiency (see e.g. Kaiserman 2016, Beebee & Kaiserman 2020). These intuitive ideas are spelled out in a variety of ways in the literature, normally in terms of conditional probabilities (see also Kaiserman 2018, Demirtas 2022). The details of the different accounts need not concern us; it will suffice here to capture the intuitive ideas.

The basic intuitive difference between degree-of-necessity and degree-of-sufficiency accounts mirrors the difference in the non-degree-theoretic causation literature between necessity-based, that is, difference-making, accounts and sufficiency-based accounts. Counterfactual analyses, for example, are necessity-based accounts: a cause is necessary for the effect in the sense that without the cause, the effect would not have occurred. Similarly, degree-of-necessity accounts of degrees of causation provide measures for *how* 'necessary' the various different causes were for the effect. To put it in more intuitive terms, these accounts propose to measure *how much* of a difference the various different causes made. Consider a forest fire (f), which starts when lightning hits a tree (l). The presence of oxygen (o) made a big difference to f compared to l: without o no fire could have started, whereas (thanks to the presence of a careless smoker) there was a decent chance that f would still have happened without l. Sufficiency-based degree-theoretic accounts, by contrast, compare how 'close' to sufficient x is for y with how close to sufficient z is for y. In the forest fire case, o was nowhere near close to sufficient for f; after all, oxygen is always present in forests and fires don't often start. l was much closer to sufficient: given l, f was very likely.

Besides these two broad approaches to specifying degrees of causation, there are at least five others.

First, Moore (2009) recommends linking degrees of causation to temporal distance: roughly, x is more of a cause of y than z is just when x is closer to y in time than is z.

Second, there is an approach based on amount of conserved quantity transferred (Demirtas 2022): x is more of a cause of y than z is just when more of a conserved quantity is transferred from x to y than from z to y.

Third, there is an approach that applies when an effect is caused in multiple, parallel ways (see Braham and van Hees 2009). Roughly, x is more of a cause of y than z is when for causings of y, c_1, c_2 ..., c_n, x is a part of more causings of y than z is.

Fourth, we can define degrees of causation modally: x is more of a cause of y than z is just when the number of worlds in which x causes y is greater than the number of worlds in which z causes y. (This approach may utilize more sophisticated measures over worlds; see Halpern and Hitchcock 2015). Relatedly, we could use probabilistic counterfactuals: if x and z together make y very probable, x is more of a cause if y would have been more probable with only x rather than only z.

Fifth, we can define degrees of causation in terms of strengths of disposition or causal power. Thus, x is more of a cause of y than z is if x's disposition to cause y or power to cause y is greater than z's (discussed in Sartorio 2020; Mumford and Anjum 2011).

And on it goes. This non-exhaustive survey of approaches teaches us that degrees of causation are generally specified in terms of something *else*, and usually in terms of some graded notion that we can specify—initially at least—independently of any attempt to link it to degrees of causation. Indeed, a feature of the world that we usually need for other purposes—such as probability, modality, counterfactual dependence and so on—is generally being co-opted into an account of degrees of causation. Another lesson comes from the diversity of views surveyed. The options for specifying degrees of causation are many, and it is not clear what, if anything, unites them.

4. Causation need not come in degrees

In this section, we use the diversity of options considered in §3 to undermine the Practice Argument. Recall the form of that argument:

The Practice Argument

(7) Practice *P* is reasonable.

(8) Practice *P* is reasonable only if causation comes in degrees.

Therefore,

(9) Causation comes in degrees.

As noted, we concede the first premise: there are a number of practices—such as the practice of apportioning different degrees of culpability for a particular outcome—that we agree are reasonable. We reject the conditional (8), that the reasonableness of these practices implies that causation comes in degrees. For this implication to hold, it must be the case that the reasonableness of the practices at issue cannot be accounted for without degrees of causation. This, we maintain, is false. For there are other graded notions that can underpin the reasonableness of the practices in question without being linked to degrees of causation. Among these notions are the very features introduced in the previous section for specifying degrees of causation.

To show this, we'll consider a couple of concrete examples and show how to account for the practice of apportioning different degrees of culpability without appealing to degrees of causation. The first example we'll consider is taken from Sartorio:

> ... each of three people (*X*, *Y*, and *Z*) released ten litres of toxic liquids in the lake. Imagine that the circumstances were such that all ten litres contributed to the lake's pollution in the case of one person, *X* (no evaporation of the toxic substances at all), zero litres in the case of a second person, *Y* (complete evaporation of the toxic substances), and only one drop in the case of a third, *Z* (close-to-complete but not fully-complete evaporation). If there are no degrees of causal contribution for culpability to track, then ... there is no difference in culpability between *X* and *Z* (Sartorio 2020: 358–9).

In this case, the claim is that degrees of culpability require degrees of causation. Without degrees of causation, the thought goes, there is no reasonable basis upon which to apportion degrees of culpability.

We disagree. There is a way to reasonably apportion culpability along broadly causal lines without appealing to degrees of causation, by appealing to a combination of non-degree-theoretic causation and some other gradable quantity or feature. In the case at hand, we can apportion different levels of culpability to *X* and *Z* straightforwardly by focusing not on degree of causation, but instead degree of pollution. *X*'s dumping caused more pollution than did *Z*'s dumping. We can thus say that *X* is more culpable for polluting the lake because he caused more pollution than did *Z*.

To see the point a bit more clearly, it is useful to compare two explanations of the different degrees of culpability borne by *X* and *Z*.

A *X* is more culpable for polluting the lake than *Z* because *X* is more of a cause of the lake's pollution than is *Z*.

B *X* is more culpable for polluting the lake than *Z* because *X* caused more of the lake's pollution than did *Z*.

What we are suggesting, then, is that degree of culpability in this case can reasonably track the amount of pollution caused by *X*, *Y* or *Z* rather than the degree to which *X*, *Y* or *Z* caused the pollution. Here we are effectively trading one graded notion—degree of causation—for another—degree of pollution. Now, our claim is *not* that the explanation in terms of degree of pollution is better, or the correct explanation. Our claim is just that it is at least as good as the degree of causation explanation when it comes to providing a basis for apportioning culpability. Culpability, in this case, can be apportioned by looking at *how much x was caused* rather than *how much one caused x*.

We expect this approach to easily generalize to other cases where there is a quantifiable effect. On plausible consequentialist approaches, for example, we are generally interested in the amount of harm that was produced. For a given harm, we can either attach culpability to the extent to which that harm was caused by an actor *or* we can attach culpability to how much harm an actor caused. Given that this second way of apportioning culpability will be available in most cases, there

should generally be a way to reasonably apportion culpability without using degrees of causation. That, however, is enough to break the implication from the reasonableness of apportioning culpability to the existence of degrees of causation.

This general solution carries over to many cases of legal responsibility as well. That's because, in many such cases, we are interested in a quantifiable effect—again, typically a harm of some kind. For instance, imagine that a company goes bankrupt, and civil proceedings begin against two directors *X* and *Y* on the grounds of negligence. Moreover, suppose that due to *X*'s actions, the company lost a million dollars, and due to *Y*'s actions the company lost a hundred thousand dollars. We want to say that *X* bears a greater degree of liability than does *Y*. One way to do this is by appealing to degrees of causation: *X* was more of a cause of the bankruptcy than was *Y*. Another option, however, is to consider how much of the drain on finances that led to bankruptcy was due to each person. Because *X*'s actions drained the finances of the company more than *Y*'s, we can apportion liability accordingly: *X* was more responsible than *Y* because *X*'s actions drained the finances of the company more than *Y*'s. Thus, we have two candidate explanations for the differing levels of legal responsibility in this case:

A′ *X* is more responsible for the bankruptcy than *Y* because *X* is more of a cause of the bankruptcy than *Y* is.

B′ *X* is more responsible for the bankruptcy than *Y* because *X* caused more of the company's financial losses than did *Y*.

As before, the point is just that B′ seems like a perfectly reasonable basis upon which to apportion legal responsibility, and so there is no reason to suppose that degrees of causation are required in the way needed to run the Practice Argument.

At this point, one might respond that these are easy cases. A more challenging case is one where the degree of responsibility seems to come apart from how much of a given effect was caused. To see this, suppose we modify the bankruptcy case. As before, suppose that due to *X*'s actions, the company lost a million dollars, and due to *Y*'s actions the company lost a hundred thousand dollars. This time, however, suppose that *X* was only mildly negligent, in a way that (let's assume) many company directors routinely are and nobody really cares unless things go very badly as a result. Suppose, by contrast, that *Y* was egregiously negligent, taking out unauthorised unsecured loans, failing to do due diligence, and so on. Now, suppose we want to work out who is more liable for the bankruptcy. If what we are interested in is *negligence*, then it seems *Y* is the most guilty party, not *X*. We cannot explain this practice of apportioning liability by appealing to how much of the finances each person drained.

Note that in this case, the effect is quantifiable. But we cannot apportion liability by focusing on how much of the effect was caused by each party. A more extreme case is one in which the effect is not quantifiable at all. For instance, suppose that *X* and *Y* cause *Z*'s death. It does not seem right to ever say that *X* caused *more* of Z's death than *Y* did. On the other hand, one might argue, it does seem sometimes right to say that *X* was more of a cause of *Z*'s death than *Y* was.

Focus on the bankruptcy case. There are two things to say here. First, note that degree of causation does not obviously do better. For it is not clear that *Y* is more of a cause of the bankruptcy than is *X*. *X*'s actions, after all, led to the company losing large sums of money. *Y*'s actions did not. The second thing to say is that *negligence comes in degrees*. Thus, what we have in this case is *another* graded notion that can be used to apportion legal responsibility. Thus, instead of saying that *Y* is more legally responsible for the bankruptcy because they were more of a cause of the law's being broken, we can say that *Y* is more legally responsible in this case because they were more negligent. Degree of negligence seems like a perfectly reasonable way to apportion legal responsibility. Moreover, it can be used when very negligent acts did not cause much of an outcome, as is the case in the bankruptcy example under consideration.

This last point helps with cases in which there is no quantifiable effect. Take the case in which *X* and *Y* cause *Z*'s death. Insofar as it seems reasonable to say that x is more of a cause of *Z*'s death than was

Y, that is likely because of some other feature of the case. Perhaps *X* bludgeoned *Z* a bit more than *Y* did. Or perhaps *X* did more of the planning than *Y* did. Whatever the story is, we can simply appeal to this other graded feature as the basis for apportioning responsibility. Thus, we can appeal to (say) bludgeoning facts directly: *X* is more responsible for *Z*'s death than *Y* is because *X* bludgeoned harder. We can even set degrees of responsibility proportional to degrees of bludgeoning, though we don't recommend this in general. A better way forward is to appeal to what the bludgeoning does to the effect: perhaps it makes it more likely to occur, or perhaps more bludgeoning means more necessity for the effect, or whatever.

This helps bring into focus a more general point. We have already shown that degree of responsibility can often be reasonably apportioned by focusing on how much of an outcome one caused, rather than how much one caused a specific outcome. The discussion of negligence and bludgeoning, however, reveals a second way to apportion responsibility without degrees of causation in those cases in which we don't or cannot say that degree of responsibility tracks how much of an outcome one caused. What we can do, instead, is look at a range of other graded notions, such as the very graded notions considered in the previous section.

For notice that each of these is something that we can specify independently of degrees of causation. Degrees of sufficiency or necessity, distance in time, amount of energy conserved, and so on, are all things we can reasonably believe in *without* accepting that there are degrees of causation. Thus, it is compatible with causation not coming in degrees that each of these phenomena is there and available to do work.

With that in mind, we can contrast two pictures. On the one hand, we can have degrees of causation plus each of the phenomena considered in the previous section, with some or all of those being linked to degrees of causation in some respect. On the other hand, we can have causation without degrees plus each of the phenomena considered in the previous section. If the Practice Argument is correct, then on this second picture there should be no reasonable way to apportion different degrees of moral or legal responsibility. We have already suggested that this is incorrect to a certain extent. But we can now see a second respect in which this is wrong: the very phenomena that we considered in the previous section can be used as the basis for reasonably apportioning moral and legal responsibility directly, without using degrees of causation.

To see this, consider a moral case. Imagine that three bullies *X*, *Y* and *Z* are teasing *S*. Imagine also that there is a hierarchy among the bullies so that neither *Y* nor *Z* will bully *S* unless goaded into it by *X*, and *Z* will not bully *S* unless goaded by at least one of *X* and *Y*. Suppose we grant that this hierarchy brings with it degrees of culpability: *X* is more culpable than *Y* who is more culpable than *Z* for the bullying. One way to account for this is to use degrees of causation linked to degrees of necessity. *X* is necessary for the bullying to the highest degree: without them, the bullying won't happen. *Y* is not necessary for the bullying, because without them *X* will still bully, but they are closer to being necessary for the bullying than is *Z*, who in turn has the lowest degree of necessity for the bullying because without them, *X* and *Y* will still bully *S*. Using these degrees of necessity, then, we can conclude that *X* is more of a cause of the bullying than *Y* or *Z*, *Y* is more of a cause than *Z*, and *Z* has the least causal efficacy. Degree of culpability can then be distributed accordingly among the bullies.

That's one way to apportion culpability. Another way, however, is to appeal directly to degrees of necessity. Thus, we can say that *X*, *Y* and *Z* are all causes of the bullying. However, they don't all bear the same degree of culpability for the bullying because they differ in their degree of necessity for the bullying, where this degree of necessity is not linked to degree of causation. This gives us two explanations:

A″ *X* is more culpable than *Y* who is more culpable than *Z* for bullying *S* because *X* was more of a cause than *Y* who was more of a cause than *Z* for the bullying.

B″ *X* is more culpable than *Y* who is more culpable than *Z* for bullying *S* because, despite all being causes of the bullying, *X* was more

necessary than Y who was more necessary than Z for the bullying.

B″ appears to be a perfectly reasonable basis upon which to apportion culpability. Indeed, to a certain extent this is conceded by our opponent. That's because the way that degrees of causation are used to explain the way that culpability is apportioned rests on the link to degrees of necessity. So, in a sense, if all sides can agree about the degrees of culpability, all sides also agree that degrees of necessity can be used to apportion moral responsibility. Our point is just that they can do that *without* being linked to degrees of causation.

This leads us to a more general argument showing that an alternative to using degrees of causation to apportion moral or legal responsibility will always be available. Suppose that we want to apportion responsibility based on degrees of causation. Thus, we say that X's degree of responsibility is higher than Y's with respect to some outcome o because X's action, x, was more of a cause of o than Y's (y) was. For any such explanation, we can ask the further question: why is it the case that x was more of a cause of o than y was? In order to answer this further question, we might appeal to one of the approaches to specifying degrees of causation considered in the previous section. So, for instance, we might say that the probability of o on x is greater than the probability of o on y. Or we might say that x is more necessary for o than y is, and so on. Indeed, some such approach will be needed to explain why x is more of a cause of o than y is. But whatever we use to explain why x is more of a cause than y is, we can use that very feature to directly explain why x is more culpable for o than y is, in whatever sense. So, for instance, if x is more of a cause of o than y is because x is closer to sufficient for o than is y, we can just use this fact to directly explain the differential apportioning of responsibility.

The strategy we are proposing is thus one of cutting out the middleman. Explanatorily speaking, degrees of causation mediate the explanation for why x is more responsible for o than y is. Nothing seems to be lost if we just remove degrees of causation and directly connect responsibility to whatever graded notion is being linked to degrees of causation. If that's right, though, then it seems there will always be available an alternative way to explain why a given practice is reasonable. We can just appeal to whatever it is that putatively explains why something is more of a cause than something else, and use that thing to apportion responsibility, without any reference to degrees of causation. Thus, the Practice Argument admits of a very general kind of response.

Note that in outlining this general response to the Practice Argument, we have focused on moral and legal responsibility. As discussed in §2, however, there are other practices that one could use to run that argument, such as the scientific practice of apportioning degrees of explanation. We haven't said much about explanation, but we see no reason why our argument shouldn't generalize to scientific practice. For example, Northcott (2008b; 2013) argues that we should apportion degrees of explanatory strength on the basis of degrees of causal strength, where the latter is understood in probabilistic terms. As above, we can simply cut out the middleman and associate degrees of explanatory strength with conditional probabilities directly. Moreover, as a number of authors have argued, there are non-causal scientific explanations. Many of these also come in degrees. That being so, a purely causal account of degrees of explanation won't do.

One might object at this point that our 'cutting out the middleman' strategy is one that applies equally to *any* broadly reductive analysis of anything. Accordingly, given that it is just not true that eliminativism is always a better option than reductionism, we need some special reason to endorse eliminativism when it comes to degree-theoretic causation. One way to see the point is to note that our strategy could be applied to causation in general. For instance, suppose one offers an analysis of causation in terms of counterfactual dependence. For the analysis to be plausible, counterfactual dependence must do everything that the notion of causation can do. But then why not just dispense with causation entirely, using counterfactual dependence instead? So far, however, we only aim to have argued that eliminativism is *a* viable option for degrees of causation, not that because we can cut out the middleman, we

must. To show that eliminativism is a better option than reductionism we need to show that cutting out the middleman is the better option; it is to that task that we now turn. Having done that, we will have achieved an appropriate contrast with the case of ordinary causation. For while cutting out the middleman in the case of ordinary causation is indeed an option, we lack analogous arguments that would force the cut.

5. Doing away with degrees

The argument of §4 involved showing that those responsibility-attributing practices which might seem to deploy degrees of causation do not *have* to be understood as doing so. But of course, it is open to the proponent of degrees of causation to respond that all this doesn't amount to an argument against degrees of causation. The mere possibility that one can do without degrees of causation doesn't provide grounds for those who do accept them to abandon them. Nor need it be persuasive to any neutral parties to the debate, who may be looking for positive grounds to favour one side.

We think there are positive grounds for doing away with degrees. In this section, we spell out those grounds by drawing some morals from our objections to the Argument from Practice in the previous section. According to the Argument from Practice, we need degrees of causation to do important work: assigning degrees of responsibility, accounting for the fact that explanations can be partial, and so on. And, as we saw in §3, there are various different candidate features of the world that different degree-theoretic accounts take to deliver degrees of causation.

Our basic idea, then, is this. A given degree-theoretic account privileges just one underlying feature (degree of necessity, say, or a conditional probability of one kind or another) and identifies the degree of that feature as the degree of causation. But what about the other features? Do *they* have any importance when it comes to assigning responsibility, or understanding partial explanation, or whatever? If they do, then we have no reason to think that *one* feature is privileged: we will have to accept that all of the various different causation-adjacent features have important work to do, and no reason to dignify one of them as the one that *is* the degree of causation. And, we'll argue, the various different features do different and important work.

Consider the example of the writing of this paper, to which all of its authors contributed in various different ways. Suppose that, while she put in almost no actual work, the guiding idea—without which the paper might not have been written at all—was *A*'s. By contrast, *B* put in most of the hours but the work was largely fungible: constructing the bibliography, chasing down references, tidying up, and so on. *C*, by contrast again, actually wrote the majority of the words you are now reading. These are all (to us) important facts about how the paper came to be. They go towards determining, for example, how much gratitude is due from each of us to our co-authors, or how much obligation there is on any individual to do whatever rewriting or reformatting or whatever needs to be done once the paper is written, and so on.

How might we go about assigning degrees of causation in this case? Well, *A* scores highest when it comes to degrees of necessity, while perhaps *B* scores highest when it comes to degrees of sufficiency. If we think instead in terms of some kind of measurable quantity, we get different answers depending on what quantity we're measuring: words? Effort? Time? Philosophical acumen? And so on. If we had to single out just one of these features to identify as the degree of causation, which one should we pick? That question, we submit, has no sensible answer. Sure, you might ask, in a given context, whether or not *X* contributed more than *Y* did. If we're lucky, the context will determine the kind of answer required. If we aren't lucky, you'll need to specify: contributed more *what*? More time? More effort? More brain power? The answer, 'more *causing*', does not help at all. The other co-authors might rightly feel the most gratitude towards *A* for having the idea that made the paper possible in the first place, while also thinking that it is probably her job to do the rewriting given how little actual work she put in. No single measure of degree of causation will capture all of the work that the degrees of causation-adjacent features

might be needed for. Our advice to a micro-managing funding body that wanted the co-authors to quantify their respective contributions would be that they can obsess about the components just described if they really must. But doing so is pretty unhelpful, as it obscures what really matters.

The same point goes for some of the examples discussed in the previous section. Take *X*, *Y* and *Z*, who were bullying poor *S*. *X* was the ring-leader, and his actions had the highest degree of necessity: had he not bullied *S*, the others most likely would not have pitched in. On the other hand, it was actually *Z* who did most of the bullying, *Y* who kept a look-out to make sure no teachers could see, and so on. Who is most liable for punishment—or, alternatively, liable for the most punishment? That's a difficult question. But, again, issuing a verdict on who was more of a cause by picking on one of the relevant features that come in degrees gets us no further in answering it than we already were.

This is a problem for degree-theoretic accounts because it requires us to accept that *any* of the various different features that come in degrees might be relevant and important for some purpose or other or in some context or other. And once that's granted, singling one of them out and identifying it with the degree of causation is not just, as we argued in §4, unnecessary; it's invidious. It suggests that that particular feature is *more* important—*simpliciter*—than the others. And no one feature *is* more important *simpliciter.*

At this point, our opponent might concede that no one feature is more important *simpliciter,* but suggest one of two potential solutions to the problem. First, they might suggest that we should simply be pluralists with respect to degrees of causation.[1] Kaiserman, for example, suggests that 'in describing *A* as "more of a cause" of some effect than *B*', there are various 'different things' we might be 'measuring', and '[c]onsidered apart from their potential applications, none of these measures is any 'better' than the others' (2018:5).

1. Thanks to an anonymous referee for suggesting this response.

The problem with this response is that it is hard to see why all the different things we might be measuring deserve to count as degrees of *causation*. As we saw in §3, degree-theoretic accounts come in many different stripes. These parallel, to at least some extent, extant non-degree-theoretic accounts of causation: necessity-based accounts, sufficiency-based accounts, causal power accounts, conserved quantity accounts, and so on. If we do not want to embrace a pluralist non-degree-theoretic account of causation and endorse some univocal account of causation *C*, then we are left with the question: if causation is analysed as per *C*, why are the remaining things we are measuring in degree-theoretic terms degree of *causation* at all?

On the other hand, we might marry pluralism about causation-by-degrees with a parallel pluralism about non-degree-theoretic causation. Then each of the 'different things' we are measuring in degree-of-causation talk corresponds to some *bona fide* species of causation. But that would constitute a much more radical pluralism about causation than the kind of pluralism that has been defended by, *inter alia*, Hall (2004), who argues that we should think of causation as coming in *two* kinds: difference-making and production. That kind of limited pluralism, however, is a far cry from the very radical pluralism being proposed here.

Second, our opponent might claim that the problem can be avoided by adopting a *multidimensional* conception of degrees of causation, according to which the various different factors are aggregated into an all-things-considered verdict about the degree to which *x* caused *y*. We disagree: *no one thing*—not even a multidimensional thing—can fulfil the role intended for degrees of causation. The argument we offer to this conclusion is the *Argument from Aggregation*.

Here is the idea. Degrees of causation must generate a stable overall ranking among causes of *e* to govern the allocation of responsibility for *e*. But there is no way of generating a stable ranking to correspond to the comparative *... is overall no less a cause of e than ....* And so, there is nothing that can play the role that degrees of causation are intended to play. Here's the argument:

The Argument from Aggregation

A1 Overall degrees of causation must yield a single aggregated comparative ranking.[2]

A2 The comparative *... is overall no less a cause of e than ...* is *multidimensional* (McConnell-Ginet 1973; Sassoon 2013), sensitive to many underlying comparative contributions to the effect *e*.

A3 Constraints on aggregation of underlying dimensions into an overall ranking, plausibly required of *... is overall no less a cause of e than ...*, are jointly unsatisfiable.

A4 So, *... is overall no less a cause of e than ...* does not express a single aggregated ranking. (From A2, A3)// *Therefore*,

C The role of degree of causation is unsatisfiable; there are no degrees of causation. (From A4, A1)

Since the argument is valid, the key issue is a defense of A1, A2, and A3. Consider A1. We should accept A1 because a notion of 'degree of causation' that does not yield a single coherent ranking would not be sufficient for allocations of responsibility, and so would not fit the role degrees of causation are supposed to play. For instance, what matters in who deserves a greater punishment for some blameworthy outcome is who is most responsible overall. If degree of causation is to be intimately linked with responsibility, as part of its conceptual role, then it must determine overall comparative rankings of causal contribution to feed into overall rankings of responsibility (where those exist). So, A1 is true.

A2 is plausible. A *dimensional* predicate, like *tall* or *expensive*, is one whose application to an object is tightly bound to the existence of a dimensional scale (such as height in the case of *tall*, or price in the case of *expensive*), and to the position (absolute or relative) of the object on that scale (Kennedy 2007). Multidimensional predicates

> are associated with multiple dimensions... An adjective like *healthy*, for example, may be associated with many different dimensions simultaneously, such as blood pressure, cholesterol, cancer, lung functions, pneumonia, chickenpox, and so forth. One can be healthy with respect to blood pressure, but not with respect to cholesterol. (Sassoon 2013: 336; cf. D'Ambrosio and Hedden 2024: §1)

As we noted earlier, following Sartorio (2020), there seem to be many rankings that are pertinent to the relative extents to which x and y causally contribute to a given e, including but not limited to the underlying dimensions of necessity and sufficiency (Demirtas 2022). Thus *... is overall no less a cause of e than ...* is, like *healthy*, a multidimensional predicate.

The final premise is A3, that the constraints on aggregating underlying dimensions into a ranking are not jointly satisfiable. It is possible to satisfy those constraints when *... is overall no less a cause of e than ...* denotes a relation that is not a weak ordering. In particular, if that relation is allowed to be incomplete, leaving some events uncompared with respect to causal contribution, then it could be the product of appropriately aggregating underlying dimensions. But, as our argument from aggregation suggests, an incomplete relation on the causes of e won't satisfy the role degrees of causation plays in attributions of responsibility.

The key to understanding and defending A3 is to note a strong mathematical and conceptual parallel between the aggregation of underlying dimensions of a multidimensional predicate, and the aggregation of individual preference rankings into a social 'betterness' ranking. The parallel can be exploited to defend A3 by appealing to Arrow's Impossibility Theorem (Arrow 1950), which shows that plausible constraints on the aggregation of individual preference rankings to generate an overall 'social' ranking *cannot* be simultaneously satisfied. These constraints are as follows (Arrow 1950; Peterson 2017: 295–7; Pettigrew 2020: §6.1):

2. Specifically, a *weak ordering*: a reflexive, complete, and transitive relation on the set of causes of some outcome.

1. The aggregation can handle any possible pattern of underlying rankings. (*Unrestricted Domain*)
2. If x is ranked ahead of y on all underlying rankings, then it should be ranked ahead on the aggregate ranking. (*Weak Pareto*)
3. The aggregate ranking of x with respect to y depends only on the values of x and y on each underlying dimension, and not on how they stand with respect to any third thing. (*Independence of Irrelevant Alternatives*)
4. The aggregate ranking isn't always determined by one underlying dimension. (*No Dictator*)

It was noted almost immediately that Arrow's result could be applied to many aggregation problems in addition to social choice.[3] Each application can be seen in turn as a manifestation of a general problem: how to construct comparative overall rankings when many underlying dimensions need to be aggregated (D'Ambrosio and Hedden 2024: 262, Grinsell 2012).

Accordingly, given a multidimensional predicate such as *healthy* and some underlying collection of rankings along various dimensions of health, if we want *overall healthier* to be a ranking that is generated by those underlying dimensions (no matter what they happen to be), then either one dimension of health must trump all the others; or whether x is healthier than y overall must depend on where x and y stand on some underlying ranking to some other thing z; or it might turn out that x is overall worse than y even though x is better than y in every respect.

The application to degrees of causation—specifically to the multidimensional comparative *... is overall no less a cause of e than ...*—is straightforward, assuming that the relevant constraints on aggregation apply to degrees of causation. We think it plausible that these apply to degrees of causation. Weak Pareto is almost platitudinous. If overall causal contribution is to sum up many respects in which events might contribute to outcomes, then it must do so in a way that is guided by those dimensions. If all underlying rankings agree on the relative contributions of x and y, then the overall ranking must do so as well, in order to count as aggregating other dimensions, rather than curating another ranking of its own.

Unrestricted Domain might seem ambitious, but (i) the proof can go through with tighter constraints on possible rankings (Gaertner 2001), and (ii) weakening this condition in practice has the effect of ensuring that some combinations of underlying rankings fall outside the scope of the aggregation function and fail to give rise to an overall ranking—so also fail to permit the overall ranking of causes to feed into allocations of responsibility, etc.

No Dictator is plausible because degrees of causation are non-trivially multidimensional and the principle is also theoretically desirable for the proponent of degrees of causation. If No Dictator is violated then we may effect a wholesale identification of degrees of causation with some particular respect in which events might cause an outcome to a greater or lesser degree.

Independence of Irrelevant Alternatives is discarded by some approaches to applications of preference aggregation, particularly in voting theory (see, e.g., Peterson 2017: 294). But it seems implausible to do so when it comes to degrees of causation. Could it plausibly be that variation in whether X is overall no less a cause of the bankruptcy than Z could depend on whether X caused more of the financial loss than some other Y, even though no causal dimension changes its view on the ranking of X and Z? Permitting such a case to arise would involve denying that the overall ranking reflects the aggregated view of the relative contributions of X and Z.

We are left with the conclusion that the most natural way to aggregate in light of Arrow's result and the plausible constraints on ranking causal contributions is to abandon the idea that the output of the aggregation procedure is a ranking (or perhaps deny there is a single aggregation procedure, which leads to similar incompleteness). The most obvious path to this outcome is to deny completeness: to deny

3. For example, multi-criterial individual preference (May 1954) or the determination of overall similarity from respects of similarity (Morreau 2010).

that every pair of causes of Z is such that one is at least as much of a cause of Z as the other (D'Ambrosio and Hedden 2024: 266). But the role standardly envisaged for degrees of causation in allocating responsibility appears to require that it be a ranking: this is fairly explicit in our earlier discussion (section 2) of the role intended for degrees of causation in the practice of attributing moral responsibility. So in general, it seems, there is no coherent notion of degrees of causation.

Now, one might worry that Arrow's result and its generalizations appear to apply to any comparatives that aggregate multiple underlying rankings and so may also apply to notions like *... is no less responsible for E than ...* or *... is more culpable for E than ...*, in which case this might seem like a perfectly general problem, not one specific to degrees of causation. We think not. There are many varieties of responsibility—legal, moral, fiduciary—and we do not generally aggregate them into some all things considered notion of responsibility.

We should stress that none of this is supposed to undermine the Argument from Entailment. If your preferred theory of causation has as a consequence that causation comes in degrees, fine. We have merely argued, in effect, that you shouldn't be touting your preferred theory's degree-theoretic nature as a *virtue* of the theory. It isn't a virtue; it's just a consequence.

References

Arrow, Kenneth J. (1950). A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4): 328–46. DOI:10.1086/256963

Beebee, Helen & Alex Kaiserman (2020). Causal contribution in war. *Journal of Applied Philosophy*, 37(3): 364–77. DOI:10.1111/japp.12341

Bernstein, Sara (2017). Causal proportions and moral responsibility. In David Shoemaker (ed.), *Oxford Studies in Agency and Responsibility*, Volume 4. Oxford: Oxford University Press, 165–82.

Braham, Matthew & Martin van Hees (2009). Degrees of causation. *Erkenntnis*, 71(3): 323–44. DOI:10.1007/s10670-009-9184-8

Collins, Darrell J. & David R. Shanks (2006). Summation in causal learning: elemental processing or configural generalization? *Quarterly Journal of Experimental Psychology*, 59(9): 1505–80. DOI:10.1080/17470210600639389

D'Ambrosio, Justin & Brian Hedden (2024). Multidimensional adjectives. *Australasian Journal of Philosophy*, 102(2): 253–77. DOI:10.1080/00048402.2023.2277923

Demirtas, Huzeyfe (2022). Causation comes in degrees. *Synthese*, 200(1): 1–17. DOI:10.1007/s11229-022-03507-2

Gaertner, Wulf (2001). *Domain Conditions in Social Choice Theory*. Cambridge: Cambridge University Press.

Gerstenberg, Tobias; Noah D. Goodman; David A. Lagnado & Joshua B. Tenenbaum (2021). A counterfactual simulation model of causal judgements for physical events. *Psychological Review*, 128(5): 936–75. DOI:10.1037/rev0000281

Goh, Joshua (2023). It's difficult to explain away the appearance that causation comes in degrees: a reply to Sartorio. *Journal of Applied Philosophy*, 40(1): 109–22. DOI:10.1111/japp.12615

Grinsell, Timothy W. (2012). Avoiding predicate whiplash: social choice theory and linguistic vagueness. *Semantics and Linguistic Theory*, 22: 424–40. DOI:10.3765/salt.v22i0.2628

Halpern, Joseph Y. & Christopher Hitchcock (2015). Graded causation and defaults. *British Journal for the Philosophy of Science*, 66(2): 413–57. DOI:10.1093/bjps/axt050

Hall, Ned (2004). Two concepts of causation. In John Collins, Ned Hall & L. A. Paul (eds.), *Causation and Counterfactuals*. Cambridge, MA: MIT Press, 225–76.

Henne, Paul; Kevin O'Neill; Paul Bello; Sangeet Khemlani & Felipe De Brigard (2021). Norms affect prospective causal judgments. *Cognitive Science: A Multidisciplinary Journal*, 45(1): e12931. DOI:10.1111/cogs.12931

Icard, Thomas F.; Jonathan F. Kominsky & Joshua Knobe (2017). Normality and actual causal strength. *Cognition*, 161: 80–93. DOI:10.1016/j.cognition.2017.01.010

Kaiserman, Alex (2016). Causal contribution. *Proceedings of the Aristotelian Society*, 116(3): 387–94. DOI:10.1093/arisoc/aow013

Kaiserman, Alex (2018). 'More of a cause': recent work on degrees of causation and responsibility. *Philosophy Compass*, 13(7): e12498. DOI:10.1111/phc3.12498

Kennedy, Chris (2007). Vagueness and grammar: the semantics of relative and absolute gradable predicates. *Linguistics and Philosophy*, 30(1): 1–45. DOI:10.1007/s10988-006-9008-0

Kominsky, Jonathan F.; Jonathan Phillips; Tobias Gerstenberg; David Lagnado & Joshua Knobe (2015). Causal superseding. *Cognition*, 137: 196–209. DOI:10.1016/j.cognition.2015.01.013

Lagnado, David & Shelley Channon (2008). Judgements of cause and blame: the effects of intentionality and foreseeability. *Cognition*, 108(3): 754–70. DOI:10.1016/j.cognition.2008.06.009

Lewis, David (2000). Causation as influence. *Journal of Philosophy*, 97(4): 182–97. DOI:10.2307/2678389

Liljeholm, Mimi & Patricia W. Cheng (2009). The influence of virtual sample size on confidence and causal-strength judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(1): 157–72. DOI:10.1037/a0013972

May, Kenneth O. (1954). Intransitivity, utility, and the aggregation of preference patterns. *Econometrica*, 22(1): 1–13. DOI:10.2307/1909827

McConnell-Ginet, Sally (1973). *Comparative Constructions in English: A Syntactic and Semantic Analysis*. PhD thesis, University of Rochester.

Moore, Michael S. (2009). *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. Oxford: Oxford University Press.

Morreau, Michael (2010). It simply does not add up: trouble with overall similarity. *Journal of Philosophy*, 107(9): 469–90. DOI:10.5840/jphil2010107931

Morris, Adam; Jonathan Phillips; Tobias Gerstenberg & Fiery Cushman (2019). Quantitative causal selection patterns in token causation. *PLoS ONE*, 14(8): e0219704. DOI:10.1371/journal.pone.0219704

Mumford, Stephen (2013). Causes for laws. *Jurisprudence*, 4(1): 109–14. DOI:10.5235/20403313.4.1.102

Mumford, Stephen & Rani Lill Anjum (2011). Spoils to the vector: how to model causes if you are a realist about powers. *The Monist*, 94(1): 54–80. DOI:10.5840/monist20119414

Northcott, Robert (2005). Pearson's wrong turning: against statistical measures of causal efficacy. *Philosophy of Science*, 72(5): 900–12. DOI:10.1086/508118

Northcott, Robert (2008a). Can ANOVA measure causal strength? *Quarterly Review of Biology*, 83(1): 47–55. DOI:10.1086/529562

Northcott, Robert (2008b). Weighted explanations in history. *Philosophy of the Social Sciences*, 38(1): 76–96. DOI:10.1177/0048393107311045

Northcott, Robert (2013). Degree of explanation. *Synthese*, 190(15): 3087–105. DOI:10.1007/s11229-012-0124-9

O'Neill, Kevin; Paul Henne; Paul Bello; John Pearson & Felipe De Brigard (2022). Confidence and gradation in causal judgement. *Cognition*, 223: 105036. DOI:10.1016/j.cognition.2022.105036

Peterson, Martin (2017). *An Introduction to Decision Theory*, 2nd edition. Cambridge: Cambridge University Press.

Pettigrew, Richard (2020). *Choosing for Changing Selves*. Oxford: Oxford University Press.

Sartorio, Carolina (2020). More of a cause? *Journal of Applied Philosophy*, 37(3): 346–63. DOI:10.1111/japp.12370

Sassoon, Galit W. (2013). A typology of multi-dimensional adjectives. *Journal of Semantics*, 30(3): 335–80. DOI:10.1093/jos/ffs012

Schlottmann, Anne & Norman H. Anderson (1993). An information integration approach to phenomenal causality. *Memory and Cognition*, 21(6): 785–801. DOI:10.3758/BF03202746

Sprenger, Jan (2018). Foundations of a probabilistic theory of causal strength. *The Philosophical Review*, 127(3): 371–98. DOI:10.1215/00318108-6718797