‘More of a Cause’: Recent Work on Degrees of Causation and Responsibility

Forthcoming in Philosophy Compass – please cite published version

Abstract

It’s often natural to compare two events by describing one as ‘more of a cause’ of some effect than the other. But what do such comparisons amount to, exactly? This paper aims to provide a guided tour of the recent literature on ‘degrees of causation’. Section 2 looks at what I call ‘dependence measures’, which arise from thinking of causes as *difference-makers*. Section 3 looks at what I call ‘production measures’, which arise from thinking of causes as *jointly sufficient* for their effects. Finally, section 4 examines the important question of whether there is any sense in which an agent is *more* responsible for an outcome in virtue of her action being more of a cause of it. I describe a puzzle that emerges from this question, first discussed by Bernstein (2017) and Sartorio (2015), and explore various strategies for resolving it.

1. Introduction

It’s often natural to compare two events by describing one as *more of a cause* of some effect than the other. A historian might describe nationalism as more of a cause of the First World War than militarism, for example; a policeman might describe the driver’s drunkenness as more of a cause of a car crash than the poor weather conditions; and a physicist might describe gravity as more of a cause of a particle’s acceleration than the presence of an electric field. Similar talk of ‘degrees of contribution’, of ‘causal potency’ or ‘causal efficacy’, and of ‘chief’, ‘main’ or ‘principal’ causes, is pervasive in the natural and social sciences, the humanities, and the law. Yet philosophers have traditionally paid little attention to these “principles of invidious discrimination” (Lewis 1973a: 559). “There is no way, based purely on causation, to identify one cause…as more important or significant than any other cause”, according to Wright (1988: 1146). It is thus “quite arbitrary to pick out [a cause] as ‘main’ or ‘secondary’” (Mackie 1965: 253); they “are equally essential, and no one of them can intelligibly be regarded as more basic than the others” (Nagel 1952: 162). “Philosophically speaking”, any one cause “has really no closer relation to the effect than any…other” (Mill 1868: 197-8).

Recent years have seen a softening of this attitude, however. Ever since Hume’s famous conflation of them in his *Enquiries*, the philosophy of causation has largely been shaped by two guiding intuitions: that causes *individually make a difference to*, and *jointly necessitate*, their effects. Many philosophical analyses of causation take one of these intuitions as their starting point.[1] Yet both also suggest ways of comparing causes, along several different dimensions. We can ask whether some event 'made a difference' to an effect; but we can also ask *how much* of a difference it made, or, if it didn't make a difference, *how close* it came to making a difference. We can ask whether some event was one of a plurality of events that were jointly sufficient for an effect; but we can also ask *how many* such pluralities the event belongs to, and how much it *contributed* to making each one jointly sufficient for the effect. Following Hall (2007), I call measures of the first kind 'dependence measures' and those of the second kind 'production measures'. I'll start, in sections 2 and 3, by comparing several different such measures from the recent literature and discussing their relative strengths and limitations, before turning in section 4 to the emerging debate over the relationship between degrees of causation and degrees of moral responsibility.

2. Dependence Measures

Suppose a charged particle is accelerating at a rate of 15ms^{-2} under the joint influence of gravity and an electric field. How much did each force contribute to the particle's acceleration? A natural way of answering this question would be to consider what the acceleration *would* have been had that force not been present. Suppose for example that, but for the presence of the gravitational field, the particle's acceleration would only have been 5ms^{-2}; whereas but for the presence of the electric field, it would have been 10ms^{-2}. Then there's a sense in which the gravitational field was 'more of a cause' of the particle's acceleration than the electric field, because the gravitational field made more of a difference.[2]

Over a series of papers, Northcott (2005a, 2005b, 2006, 2008a, 2008b) develops a measure of 'causal strength' according to which "the strength of a cause is how much difference it makes" (Northcott 2013: 3090). Let X and Y be variables,[3] let x_A and y_A be the actual values of X and Y respectively, and let x_C be some salient counterfactual value of X. Finally, let y_C be the value Y *would* take were the value of X to be changed to x_C by means of an intervention with

[1] For some classic examples, see Lewis (1973a) and Mackie (1965), respectively.

[2] This example is adapted from one in Sober (1988). Sober argues that "the contribution a cause makes and the difference it makes [are] one and the same issue" (Sober 1988: 303), but *only* in cases, like the charged particle case, where the causes are "commensurable in the way they produce their effects" (Sober 1988: 312). See Northcott (2005a) for a critique of this 'commensurability' requirement.

[3] Variables represent "properties or magnitudes that, as the name implies, are capable of taking more than one value" (Woodward 2003: 39). In the simplest case a variable X has two possible values – say, 0 and 1 – such that X=1 if a particular event occurs and X=0 otherwise. But multi-valued variables are possible too – for example, we might consider a variable X that can take any positive number x as a value such that X=x if and only if a certain particle is accelerating at a rate of $x$ms^{-2}.

respect to Y.[4] Now let's define the *difference made to Y* by X's taking the value x_A rather than x_C – or $\alpha([x_A, x_C], Y)$ for short – as follows:

$$\alpha([x_A, x_C], Y) = y_A - y_C \quad (1)$$

[5]

For example, the difference made to the particle's acceleration by the gravitational field's having strength 10Nkg^{-1} rather than 0Nkg^{-1} is 15 - 5 = 10ms^{-2}.[6]

In the accelerating particle case, the total effect is equal to the sum of the differences made by each cause. But this needn't be true in general. Imagine a plant that has grown to a height of 100cm with regular watering and exposure to sunlight. Had it not been watered, it would only have reached 20cm; had it not been exposed to sunlight, it would only have reached 10cm. The sum of the differences made by each cause in this case is greater than the actual height of the plant. Unlike in the accelerating particle case, then, we can't 'divide up' the plant's height in such a way that the sunlight is causally responsible for one part and the watering another.

In some cases, in fact, one or more of the causes is necessary for there being any 'quantity' of the effect at all. Imagine a primordial soup containing two chemicals that will synthesize into some organic compound given a certain activation energy.[7] Two thunderclouds pass overhead, a large one and a small one. Suppose a lightning bolt from the large cloud is more energetic than one from the small cloud, but still below the activation energy required for the chemicals in the soup to react. The combined energy of both bolts, however, is above the activation threshold. Both lightning bolts strike the soup simultaneously, and the chemicals react. The two causes made exactly the same difference to their effect in this case, namely the difference between a reaction occurring and no reaction occurring. Insofar as there is some sense in

[4] An intervention is a special kind of causal process designed to control for confounding influences and rule out 'backtracking'; Woodward (2003: 98) provides a formal definition, but the details needn't concern us here.

[5] Note that α is explicitly *contrastive* in structure – a cause 'makes a difference' to its effect only relative to a particular contrast. See Schaffer (2005), Northcott (2008c), and Kaiserman (2017a) for discussion of the alleged contrastive structure of causation.

[6] Northcott defends α against various alternative ways of measuring the relative importance of different causes, including the statistical techniques (e.g. ANOVA) typically used in the special sciences (Northcott 2005b, 2006, 2008a) and the limiting/selecting, underlying/proximate and necessary/non-necessary distinctions typically appealed to in the philosophy of history (Northcott 2008b).

[7] This case is borrowed from Northcott (2005a).

which the lightning bolt from the bigger cloud was 'more of a cause' of the reaction than that from the smaller cloud, then, α is incapable of capturing it.[8]

Another issue with α has to do with the logic of counterfactuals. Suppose Bob develops lung cancer after smoking ten cigarettes a day for thirty years. Would Bob still have developed lung cancer if he had smoked just five cigarettes a day? On some theories of counterfactuals,[9] there need be no determinate answer to this question. The closest possible worlds where Bob smokes five cigarettes a day might include worlds where he develops lung cancer and worlds where he doesn't; not (or not necessarily) because the underlying laws of nature are indeterministic, but rather because there are several different ways the condition of Bob smoking five cigarettes a day could have been realised, each of which is equally similar to what Bob actually does.

The natural solution to this problem is to move to a probabilistic framework. Let 'P(p)' denote the objective probability of p,[10] let 'do(X=x, Y)' denote the proposition that X is set to value x by means of an intervention with respect to Y, and let b be the conjunction of all salient 'background conditions'.[11] Then we can define the difference made by X taking value x_A rather than x_C to the *probability* of Y=y_A – or $\beta([x_A, x_C], Y)$ for short – as follows:

$$\beta([x_A, x_C], Y) = P(Y{=}y_A \mid do(X{=}x_A, Y) \,\&\, b) - P(Y{=}y_A \mid do(X{=}x_C, Y) \,\&\, b) \quad (2)$$ [12]

$\beta([x_A, x_C], Y)$ measures the difference made, not to Y, but to the probability of Y taking the value it actually did. For example, suppose the probability (in the circumstances) of Bob

---

8 According to Northcott, the source of the intuition that the bolt from the larger cloud is 'more of a cause' is the fact that it *would* have made more of a difference had the effect been, say, a voltage induced in a wire. "But when considering our particular effect of triggering the chemical reaction, because of the activation energy threshold I do not see how assigning different efficacies can be justified. In our *particular* example, that is, the comparison of causal efficacies must surely be trivial…even though their comparison is not trivial in *other* examples" (Northcott 2005a: 16). See also Sartorio (ms) on this point.

9 Specifically, any theory which rejects *counterfactual excluded middle*: $\forall p\ \forall q\ (p \mathbin{\Box\!\!\rightarrow} q \lor p \mathbin{\Box\!\!\rightarrow} \neg q)$. On the standard possible-worlds semantics of counterfactuals (Lewis 1973b), this is equivalent to the claim that for all p, there is at most one closest possible world in which p is true. Lewis (1973b: 80) swiftly rejects counterfactual excluded middle with the following argument: "It is not the case that if Bizet and Verdi were compatriots, Bizet would be Italian; and it is not the case that if Bizet and Verdi were compatriots, Bizet would not be Italian". For discussion and a response, see Stalnaker (1981).

10 I'm assuming here that there are non-trivial objective probabilities, even in a deterministic universe; but see Lewis (1980).

11 I assume throughout this paper that there is a difference between the causes of an effect and the "background against which the causing goes on" (Mackie 1974: 63); or, as Plato once put it, between "the real cause [and] that without which the cause would not be able to act as a cause" (*Phaedo 99a-b*). I won't speculate here on how exactly this distinction is to be drawn; but see Kaiserman (2017c).

12 Sprenger (forthcoming) proves a representation theorem to the effect that any measure satisfying certain formal constraints is ordinarily equivalent to β. Fitelson and Hitchcock (2011) call β the 'Eells measure' after Eells (1991). Versions of the same measure are also defended by Cheng (1997) and Novick and Cheng (2004).

developing lung cancer conditional on him smoking ten cigarettes a day for thirty years is 0.4, whereas conditional on him smoking just five cigarettes a day the probability is 0.1. Then Bob's smoking ten rather than five cigarettes a day made a difference of 0.3 to his probability of developing lung cancer.[13]

Some causes determinately make no difference at all to their effects. Suppose a prisoner is killed by a firing squad. Each shooter made exactly the same difference to whether (and to how) the prisoner's death occurred, namely none – had any one of the shooters failed to shoot, the prisoner would still have died, more or less exactly in the way she actually did. But effects can be *more or less* overdetermined by their causes. Chockler and Halpern (2004), for example, propose the following measure of 'causal responsibility' (see also Halpern (2016: ch.6)). Let N be the minimum number of changes that would have to be made to the other causes of Y in order to obtain a contingency where $\alpha([x_A, x_C], Y)$ is non-zero.[14] Then *how close* X's taking value x_A rather than x_C came to making a difference to Y – or $\gamma([x_A, x_C], Y)$ for short – is defined as follows:

$$\gamma([x_A, x_C], Y) = \frac{1}{1+\mathrm{N}} \qquad (3)$$

To illustrate, suppose Cathy is one of nine friends voting on whether to go to the pub or the park. Everyone (of course) votes for the pub. Cathy's vote made no difference to the outcome. But had four of her friends voted differently, Cathy's vote *would* have made a difference. So on a natural choice of variables, N in this case is 4, and $\gamma = \frac{1}{5}$. Had going to the pub won 8-1 instead, γ would have been $\frac{1}{4}$; and so on. Generally speaking the more overdetermined the cause, the lower the value of γ.[15]

[13] Note that, as Fitelson and Hitchcock (2011) point out, β exhibits 'floor effects' – $\beta([x_A, x_C], Y)$ is inversely proportional to $P(Y{=}y_A)$. In other words, if Y was very likely in the circumstances to take value y_A *whatever* happened, the maximum value of $\beta([x_A, x_C], Y)$ will be limited, since there's only so much difference X can make to the probability of $Y{=}y_A$ if it's already close to 1. Thus although β can be used to compare different causes of the same effect, it can't sensibly be used to compare causes of different effects.

[14] By 'change' here, Chockler and Halpern mean a difference of 1 in the value of an individual variable. So for example, if $X{=}2$ and $Y{=}1$, the number of changes needed to obtain a contingency in which $X{=}Y{=}0$ is 3. This means, of course, that N will depend on the variables we choose to use to model the situation of interest. For more on the issues associated with variable choice, see Franklin-Hall (2016).

[15] Hence, in contrast to α, γ *only* delivers interesting results in cases of redundant causation – in most normal cases, N=0, and γ=1.

3. Production Measures

The measures discussed in the previous section belong to a tradition that attempts to understand causation in terms of counterfactuals. Running in parallel with this tradition, however, is one that attempts instead to understand causation in terms of *sufficiency*.[16] The basic idea is that causes are *minimally jointly sufficient in the circumstances* for their effects. More precisely:

> $X_1=x_1,\ldots, X_n=x_n$ (collectively) caused $Y=y$ if and only if:
>
> i) $X_1=x_1,\ldots, X_n=x_n$ were jointly sufficient in the circumstances for $Y=y$, and
>
> ii) no proper sub-plurality of $X_1=x_1,\ldots, X_n=x_n$ was jointly sufficient in the circumstances for $Y=y$.[17] (4)

To be *a cause* of an effect is then just to be one of a plurality of events that collectively caused it (just as to be *an author* of a book is to be one of a plurality of people who collectively authored it).

Braham and van Hees (2009) develop their own measure of 'causal contribution' within this framework.[18] Let's call a plurality of events a *causing* of $Y=y$ if and only if they collectively caused $Y=y$. Now let '$\#(X_i, Y)$' denote the number of different causings of $Y=y$ to which $X_i=x_i$ belongs. Then where $X_1=x_1,\ldots, X_n=x_n$ are all the causes of $Y=y$ in our model, what I'll call the *causal power*[19] of $X_i=x_i$ with respect to $Y=y$ – or $\delta(X_i, Y)$ for short – is defined as follows:

$$\delta(X_i, Y) = \frac{\#(X_i,Y)}{\sum_{j=1}^{n}\#(X_j,Y)} \tag{5}$$

To illustrate, consider the following case from Kaiserman (2016):

[16] In particular, I have in mind accounts of causes as 'INUS' or 'NESS' conditions – see Mackie (1965) and Wright (1985) respectively. To avoid obvious counterexamples, 'sufficient' here should be understood as denoting a relation richer than mere metaphysical necessitation; see Strevens (2007) on this point. On the relationship between the sufficiency and counterfactual approaches to analysing causation, see Kment (2010).

[17] Mackie and Wright state their view in terms of 'sets of conditions' – I prefer to talk of pluralities of events, but this choice won't be relevant for what follows

[18] What follows is a dramatically simplified version of Braham and van Hees's view. For ease of exposition I skip over several interesting details, including the difference between 'minimally' and 'critically' sufficient pluralities.

[19] This is not Braham and van Hees's term, but it strikes me as an appropriate one given the close relationship between δ and Holler's (1982) 'Public Good Index' of voting power.

Committee: D_1, D_2, D_3 and D_4 are the members of an executive committee of a manufacturing company. Every committee member has one vote each, except D_1, the chair of the committee, who has two votes. The committee members all vote not to replace the company's outdated equipment. The equipment later malfunctions, injuring an employee.

Although none of the individual committee members made any difference to the outcome of the vote, there still seems to be some sense in which D_1's action was 'more of a cause' of the employee's injury than D_2's action, because D_1 had two votes and D_2 only had one. δ can capture this intuition. Let V_1, V_2, V_3 and V_4 be variables such that V_i=1 if D_i votes in favour of the motion and 0 otherwise, and let I=1 if the injury occurs and 0 otherwise. There are a total of five votes, so the minimum number of votes needed for a majority is three. Hence there are four pluralities of events that were minimally jointly sufficient (in the circumstances) for the injury: [V_1=1, V_2=1], [V_1=1, V_3=1], [V_1=1, V_4=1] and [V_2=1, V_3=1, V_4=1]. The injury was therefore caused *four times*, by four different (though overlapping) pluralities of events. Since D_1's vote contributes to three of these causings and the other votes to two causings each, we can calculate the causal power of each vote with respect to the injury as follows:

$$\delta(V_1, I) = \frac{\#(V_1, I)}{\sum_{j=1}^{4} \#(V_j, I)} = \frac{3}{2+2+2+3} = \frac{1}{3} \tag{6}$$

$$\delta(V_2, I) = \delta(V_3, I) = \delta(V_4, I) = \frac{2}{9} \tag{7}$$

As expected, then, the causal power of D_1's vote with respect to the injury is greater than that of the other votes.

Like some of the other measures considered so far, δ only delivers interesting results in certain special cases, namely those, like **Committee**, in which the effect is caused by multiple overlapping pluralities of events. In most ordinary cases, δ simply assigns each cause the same causal power. Suppose for example that a car crash was collectively caused (in the circumstances) by the driver's drunkenness and a rainstorm. Since both causes contributed to a single causing of the car crash, δ for each cause is just ½. If there is any sense in which the drunkenness was more of a cause of the crash than the rainstorm, then, δ is incapable of capturing it.

In light of this, Kaiserman (2016) proposes (what amounts to) a probabilistic extension of δ. Let D=1 if the driver is drunk and 0 otherwise, R=1 if the rainstorm occurs and 0 otherwise and C=1 if the car crash occurs and 0 otherwise. Suppose the crash was jointly caused by the drunkenness and the rainstorm, so that D=1 and R=1 were jointly sufficient in the circumstances for C=1 although neither was individually sufficient. But suppose also that $P(C{=}1 \mid do(D{=}1, C) \,\&\, b) > P(C{=}1 \mid do(R{=}1, C) \,\&\, b)$ – the probability in the circumstances of the crash occurring is greater conditional on the driver's being drunk than it is conditional on the rainstorm occurring. Then there's a sense in which the drunkenness *contributed more* to this causing of the crash, because it came closer to being sufficient for the crash by itself. More generally, where $X_1=x_1,\ldots, X_n=x_n$ collectively caused $Y=y$, we can define $X_i=x_i$'s *degree of contribution* to the causing of $Y=y$ by $X_1=x_1,\ldots, X_n=x_n$ – or $\varepsilon(X_i, [X_1,\ldots, X_n]\rightarrow Y)$ for short – as follows:

$$\varepsilon(X_i, [X_1,\ldots, X_n]\rightarrow Y) = \frac{\mathrm{P}(Y=y \mid \mathrm{do}(X_i=x_i,Y)\ \&\ b)}{\sum_{j=1}^{n}\mathrm{P}(Y=y \mid \mathrm{do}(X_j=x_j,Y)\ \&\ b)}$$ [20, 21]

To illustrate, suppose that the driver is *really* drunk, so that any number of potential distractions would have been enough for him to lose control of his vehicle. The rainstorm, meanwhile, was fairly mundane, and only contributed to the causing of the crash by slightly impeding the driver's vision. On these facts, $\mathrm{P}(C=1 \mid \mathrm{do}(D=1, C)\ \&\ b)$ is close to 1: conditional on the driver being in his inebriated state, it was very likely in the circumstances that the crash would have occurred one way or the other. On the other hand, since the rainstorm would have posed no danger to a sober driver, $\mathrm{P}(C=1 \mid \mathrm{do}(R=1, C)\ \&\ b)$ is not much greater than $\mathrm{P}(D=1\ \&\ b)$, the unconditional probability in the circumstances of the driver being as drunk as he was. Suppose for example that $\mathrm{P}(C=1 \mid \mathrm{do}(D=1, C)\ \&\ b) = 0.9$ and $\mathrm{P}(C=1 \mid \mathrm{do}(R=1, C)\ \&\ b) = 0.3$; then it follows that the drunkenness contributed to degree 0.75 to the causing of the crash and the rainstorm to degree 0.25.

Note that ε measures an event's degree of contribution *to a causing* of an effect, not to the effect itself. In particular, when an event contributes to multiple causings of an effect, ε can be used to calculate its degree of contribution to each of these causings. In **Committee**, for example, D_2's vote contributes to two causings of the injury, the causing by $[V_1=1, V_2=1]$ and the causing by $[V_2=1, V_3=1, V_4=1]$. Assuming for the sake of argument that each committee member has a probability of 0.5 of voting in favour of the motion, regardless of the actions of the others, Kaiserman (2016: 392) shows that V_2 contributes to degree 0.42 to the first causing and to degree ⅓ to the second. In just the same way that a plank of wood at the intersection of two wooden fences of different lengths can contribute to different degrees to two different 'surroundings' of the very same tree, a cause may contribute to different degrees to two different causings of the very same effect.

---

[20] The denominator of this fraction is a normalising factor included to ensure that degrees of contribution to causings always sum to 1 (i.e. $\sum_{i=1}^{n}\varepsilon(X_i, [X_1,\ldots,X_n]\rightarrow Y) = 1$).

[21] Interestingly, Rizzo and Arnold (1980) propose a theory of 'causal apportionment' in the context of the economic theory of tort law which bears some resemblance to ε, although their framework is only designed to handle comparisons between two causes.

4. Dimensions of Responsibility

Let's take stock. We've seen that there are several different ways of comparing causes – in describing *A* as 'more of a cause' of some effect than *B*, we might be saying that *A* made more of a difference (either to the effect or to its probability), that it came closer to making a difference, that it contributed to more causings of the effect, that it contributed to these causings to a larger degree, or indeed something else entirely. Considered apart from their potential applications, none of these measures is any 'better' than the others – they are simply measuring different things.

Interesting questions begin to arise, however, when we consider which, if any, of these measures should be taken into account when determining the extent of an agent's *moral responsibility* for an outcome. Most people agree that causation is necessary for moral responsibility – one cannot be held responsible for something one did not cause.[22] But is there any sense in which an agent is *more* responsible for an outcome in virtue of her actions being 'more of a cause' of that outcome?

Let's start by considering the following pair of cases, adapted from Bernstein (2017):

> **Victim**: Two independently employed assassins, unaware of each other, are dispatched to eliminate Victim. Being struck by one bullet is sufficient to kill Victim. Each assassin shoots, and the bullets strike Victim at exactly the same time. Victim dies.
>
> **Hardy Victim**: Two independently employed assassins, unaware of each other, are dispatched to eliminate Victim. Unbeknownst to both assassins, Victim is particularly hardy, and requires two bullets for his demise. Each assassin shoots, and the bullets strike Victim at exactly the same time. Victim dies.

Who is more responsible for Victim's death – the assassins in **Victim** or the assassins in **Hardy Victim**? On the one hand, it seems the assassins in **Victim** should be more responsible. After all, each assassin's shot *individually* caused Victim's death in **Victim**, whereas in **Hardy Victim** the two shots only *jointly* caused the death. Intuitively, then, while the assassins in **Victim** are both 'fully' responsible for the death, each assassin is only 'partially' responsible for the death in **Hardy Victim**.[23] Call this the 'production intuition'.

According to Bernstein, however, there's a different way of thinking about these cases. The assassins in **Hardy Victim** made a difference to whether Victim died – their shots were both

[22] Although see Sartorio (2004).

[23] Which is not to say, of course, that they are responsible for a *part* of the death – see Kaiserman (2017b: 4).

necessary, in the circumstances, for bringing about the death. In **Victim**, meanwhile, neither shot made a difference, because had any one assassin failed to shoot, Victim would still have died. Such considerations, Bernstein claims, motivate the opposite conclusion: that the assassins in **Victim** are *less* responsible than those in **Hardy Victim**. Call this the 'dependence intuition'.

Those who endorse the production intuition tend to find a close connection between an agent's degree of responsibility for some outcome and the degree of contribution her action made to the causing of the outcome (see Kaiserman 2017b). Those who endorse the dependence intuition, meanwhile, tend to find a close connection between an agent's degree of responsibility for some outcome and how close her action came to making a difference to it (see Chockler and Halpern 2004; Halpern 2016: ch.6). What we've seen is that these different approaches produce conflicting results in cases like **Victim** and **Hardy Victim**. There is thus "a puzzle about the relationship between degrees of causation and degrees of moral responsibility" (Bernstein 2017: 165).

One obvious strategy for resolving this puzzle would be to deny (and, ideally, explain away) one of the intuitions in favour of the other.[24] Alternatively, one could follow Tadros (2018) in rejecting both the production and the dependence intuitions. Bernstein's own reaction is to call for more work to be done on the metaphysical foundations of 'degrees of causation' talk. Sartorio (ms), meanwhile, argues that the best solution to the puzzle is to reject the very coherence of talk of 'degrees of causation' in the first place.

There is at least one other possible solution to Bernstein's puzzle, however: one could simply accept that there are *two* incommensurable causal dimensions to moral responsibility, the production dimension and the dependence dimension. *All other things being equal*, those who contribute to a larger degree to a causing of some outcome are more responsible for that outcome; and *all other things being equal*, those who come closer to making a difference to some outcome are also more responsible for that outcome. But there is simply no answer to the question of who is more morally responsible in cases, like **Victim** and **Hardy Victim**, which differ in opposing directions along both dimensions.

[24] I have my suspicions about the dependence intuition, for what it's worth. Suppose one of the assassins in **Victim** comes to regret what she has done. It may bring her some comfort to learn of the other assassin's existence – she may even feel a sense of relief that her action didn't make a difference to how things ultimately turned out. But I'm not sure she ought to feel any less *responsible* for the death. As a matter of fact, her action caused Victim's death – that this death was *also* caused by someone else doesn't seem, at least to me, to be relevant to her own degree of responsibility for it. As Zimmerman (1985: 115) has argued, there are good reasons to resist the view that responsibility is always "diminished simply by virtue of the fact that others...are responsible".

In defence of such a view, it's worth pointing out that similar puzzles plausibly arise for the other dimensions of moral responsibility too. Consider the following pair of cases, for example:

> **Kleptomaniac**: Annie has mild kleptomania. This manifests in the form of a recurrent desire to steal. Annie can resist this desire, but only with a non-trivial amount of effort. One day, at the airport, Annie gives in to the desire and intentionally steals Zara's suitcase from the baggage reclaim area.
>
> **Careless Thief**: Bertie is at the baggage reclaim area of the airport. He picks up what he believes to be his suitcase and walks away with it. In fact he has picked up Zara's suitcase by mistake. Bertie would have noticed this had he given the suitcase more than just a casual glance.

There is no relevant difference in the *causal* structures of these two cases. Yet we still face a puzzle when we ask who is more morally responsible for Zara's loss. On the one hand, Annie, being a mild kleptomaniac, was less sensitive to reasons not to take the suitcase than Bertie was; and all other thing being equal, those who are less sensitive to reasons are less responsible for what they do.[25] But on the other hand, Annie intentionally took Zara's suitcase, whereas Bertie didn't even know that's what he was doing (even if there's a sense in which he *ought* to have known); and all other things being equal, those who cause loss intentionally are more responsible than those who merely do so carelessly or negligently.[26] The problem is that all other things are *not* equal in **Kleptomaniac** and **Careless Thief**, and it's not clear how to weigh the different considerations against one another. Indeed, it's not clear whether it even makes sense to ask whether Annie or Bertie is more responsible – perhaps the most we can sensibly say is that they differ in opposing directions along two different dimensions of moral responsibility.

Regardless of one's views about the moral relevance of 'degrees of causation', then, it seems plausible that the *more-responsible-than* relation is not a total order – there are pairs of agents who are simply incomparable in terms of their moral responsibility for some outcome. And this presents a more general challenge about how to proceed in cases where we are called upon to make such comparisons. To make this vivid, suppose that Cathy, Duncan, and Ella comprise the cabinet of country X. They all culpably vote to declare war on country Y, which will foreseeably result in the deaths of many of Y's citizens. We can protect country Y against

---

[25] For a defence of this claim, and for different accounts of reasons-sensitivity, see Fischer and Ravizza (1998), Fischer (2012), McKenna (2013) and Sartorio (2016).

[26] Most legal systems distinguish different 'grades' of culpability, starting with intentional actions, and progressing down the scale of severity through knowledge, recklessness, and negligence ('gross' and 'ordinary'); see Simester et al. (2013: ch.5). For a classic defence of the view that we are criminally responsible for the consequences of our negligence, see Hart (1968: ch.6). For a good introduction to the recent philosophical work on the epistemic dimension to moral responsibility, see Talbert (2016: ch.5) and the papers in Robichaud and Wieland (2017).

this unjustified threat, but only by killing one of the cabinet members (this, let us suppose, will scare the others into agreeing to a ceasefire). Suppose we are morally required to pursue such defensive action. Which of Cathy, Duncan, and Ella should we target? A natural answer is that we should kill whoever is most responsible for the unjustified threat. But if the three cabinet members differ in opposing directions along different dimensions of moral responsibility, there may not be an answer to the question of who is most responsible. Although the decision-theoretic issues arising from moral incomparability are now well-known (see Hsieh 2016), this particular manifestation of the problem remains underexplored.

Conclusion

There has been a surge in philosophical engagement with the concept of 'degrees of causation' in recent years. Relatively independent debates in the philosophy of science, legal philosophy and metaphysics have led to the development of several measures of the relative importance of different causes. Given the big differences in formalism, terminology, and starting assumptions across these debates, however, it can often be difficult to determine how these different measures relate to each other, whether they are compatible, and how they fit into the wider literature on causation. The goal of this paper has been to provide a more accessible guide through this emerging conceptual landscape. As the previous section demonstrates, reflection on the connections between these measures of causation and our concepts of moral responsibility, scientific explanation (Northcott 2013), legal liability (Kaiserman 2017b) or voting power (Braham and van Hees 2009) have the potential to raise new philosophical problems and generate new avenues for theoretical enquiry.

References

Bernstein, Sara J. 'Causal Proportions and Moral Responsibility.' *Oxford Studies in Agency and Responsibility (Vol. 4)*. Ed. David Shoemaker. Oxford: Oxford University Press, 2017.

Braham, Matthew and van Hees, Martin. 'Degrees of Causation.' *Erkenntnis* 71.3 (2009): 323-44.

Cheng, Patricia W. 'From Covariation to Causation: A Causal Power Theory.' *Psychological Review* 104.2 (1997): 367-405.

Chockler, Hana and Halpern, Joseph Y. 'Responsibility and Blame: A Structural-Model Approach.' *Journal of Artificial Intelligence Research* 22 (2004): 93-115.

Eells, Ellery. *Probabilistic Causality*. Cambridge: Cambridge University Press, 1991.

Fischer, John Martin. *Deep Control: Essays on Free Will and Value*. New York: Oxford University Press, 2012.

Fischer, John Martin and Ravizza, Mark. *Responsibility and Control*. Cambridge: Cambridge University Press, 1998.

Fitelson, Branden and Hitchcock, Christopher. 'Probabilistic Measures of Causal Strength.' *Causality in the Sciences*. Ed. Phyllis McKay Illari, Federica Russo and Jon Williamson. Oxford: Oxford University Press, 2011. 600-27.

Franklin-Hall, Laura. 'High-Level Explanation and the Interventionist's 'Variables Problem'.' *British Journal for the Philosophy of Science* 67.2 (2016): 533-77.

Hall, Ned. 'Two Concepts of Causation.' *Causation and Counterfactuals*. Ed. John Collins, Ned Hall, and L. A. Paul. Cambridge, MA: The MIT Press, 2004. 225-76.

Halpern, Joseph Y. *Actual Causality*. Cambridge, MA: The MIT Press, 2016.

Hart, H. L. A. *Punishment and Responsibility: Essays in the Philosophy of Law*. New York: Oxford University Press, 1968.

Holler, Mandred J. 'Forming Coalitions and Measuring Voting Power. *Political Studies* 30.2 (1982): 262-71.

Hsieh, Nien-hê. 'Incommensurable Values.' *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition). Ed. Edward N. Zalta. URL = <https://plato.stanford.edu/archives/spr2016/entries/value-incommensurable/>.

Kaiserman, Alex. 'Causal Contribution.' *Proceedings of the Aristotelian Society* 116.3 (2016): 387-94.

Kaiserman, Alex. 'Causes and Counterparts.' *Australasian Journal of Philosophy* 95.1 (2017a): 17-28.

Kaiserman, Alex. 'Partial Liability.' *Legal Theory* 23.1 (2017b): 1-26.

Kaiserman, Alex. 'Necessary Connections in Context.' *Erkenntnis* 82.1 (2017c): 45-64.

Kment, Boris. 'Causation: Determination and Difference-Making.' *Noûs* 44.1 (2010): 80-111.

Lewis, David. 'Causation.' *Journal of Philosophy* 70.17 (1973a): 556-67.

Lewis, David. *Counterfactuals*. Oxford: Oxford University Press, 1973b.

Lewis, David. 'A Subjectivist's Guide to Objective Chance.' *Studies in Inductive Logic and Probability*. Ed. Richard C. Jeffrey. Berkeley: University of California Press, 1980. 83-132.

Mackie, J. L. 'Causes and Conditions.' *American Philosophical Quarterly* 2.4 (1965): 245-64.

Mackie, J. L. *The Cement of the Universe: A Study of Causation*. Oxford: Oxford University Press, 1974.

McKenna, Michael. 'Reasons-Responsiveness, Agents and Mechanisms.' *Oxford Studies in Agency and Responsibility (Vol. 1)*. Ed. David Shoemaker. Oxford: Oxford University Press, 2013. 151-83.

Mill, J. S. *A System of Logic, Ratiocinative and Inductive (Vol. 1)*. London: Longmans, Green, Reader, and Dyer, 1868.

Nagel, Ernest. 'Some Issues in the Logic of Historical Analysis.' *The Scientific Monthly* 74.3 (1952): 162-9.

Northcott, Robert. 'Comparing Apples with Oranges.' *Analysis* 65.1 (2005a): 12-18.

Northcott, Robert. 'Pearson's Wrong Turning: Against Statistical Measures of Causal Efficacy.' *Philosophy of Science* 72.5 (2005b): 900-12.

Northcott, Robert. 'Causal Efficacy and the Analysis of Variance.' *Biology and Philosophy* 21.2 (2006): 253-76.

Northcott, Robert. 'Can ANOVA Measure Causal Strength?' *Quarterly Review of Biology* 83.1 (2008a): 47-55.

Northcott, Robert. 'Weighted Explanations in History.' *Philosophy of the Social Sciences* 38.1 (2008b): 76-96.

Northcott, Robert. 'Causation and Contrast Classes.' *Philosophical Studies* 39.1 (2008c): 111-23.

Northcott, Robert. 'Degree of Explanation.' *Synthese* 190.15 (2013): 3087-105.

Novick, Laura R. and Cheng, Patricia W. 'Assessing Interactive Causal Influence.' *Psychological Review* 111.2 (2004): 455-85.

Robichaud, Philip and Wieland, Jan Willem. *Responsibility: The Epistemic Condition*. Oxford: Oxford University Press, 2017.

Rizzo, Mario J. and Arnold, Frank S. 'Causal Apportionment in the Law of Torts: An Economic Theory.' *Columbia Law Review* 80.7 (1980): 1399-1429.

Sartorio, Carolina. 'How to be Responsible for Something Without Causing it.' *Philosophical Perspectives* 18.1 (2004): 315-336.

Sartorio, Carolina. 'A New Form of Moral Luck?' *Agency, Freedom and Moral Responsibility*. Ed. Andrei Bcukareff, Carols Moya and Sergi Rosell. Palgrave Macmillan, 2015. 134-52.

Sartorio, Carolina. *Causation and Free Will*. Oxford: Oxford University Press, 2016.

Sartorio, Carolina. 'More of a Cause?' Unpublished manuscript.

Schaffer, Jonathan. 'Contrastive Causation.' *The Philosophical Review* 114.3 (2005): 327-58.

Simester, Andrew P. *et al. Simester and Sullivan's Criminal Law: Theory and Doctrine*. Oxford: Hart Publishing, 2013.

Sober, Elliott. 'Apportioning Causal Responsibility.' *Journal of Philosophy* 85.6 (1988): 303-18.

Sprenger, Jan. 'Foundations for a Probabilistic Theory of Causal Strength. *The Philosophical Review* (forthcoming).

Stalnaker, Robert C. 'A Defense of Conditional Excluded Middle.' *Ifs*. Ed. William Harper, Robert C. Stalnaker and Glenn Pearce. Dordrechtht: D. Reidel, 1981. 87-104.

Strevens, Michael. 'Mackie Remixed.' *Causation and Explanation*. Ed. Joseph Keim Campbell, Michael O'Rourke and Harry S. Silverstein. Cambridge, MA: The MIT Press, 2007. 4-93.

Tadros, Victor. 'Causal Contributions and Liability.' *Ethics* 128.2 (2018): 402-31.

Talbert, Matthew. *Moral Responsibility: An Introduction*. Cambridge: Polity, 2016.

Woodward, James. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press, 2003.

Wright, Richard W. 'Allocating Liability Among Multiple Responsible Causes: A Principled Defense of Joint and Several Liability for Actual Harm and Risk Exposure.' *University of California at Davis Law Review* 21.4 (1988): 1141-211.

Zimmerman, Michael J. 'Sharing Responsibility.' *American Philosophical Quarterly* 22.2 (1985): 115-22.