'A Variety of Causes', by Paul Noordhof

REVIEW BY J. DMITRI GALLOW[1]

[1] ✉: dmitri.gallow@gmail.com
🏠: http://jdmitrigallow.com

Draft of April 19, 2022

Paul Noordhof's *A Variety of Causes* presents and defends a counterfactual theory of causation. The book is *incredibly* detailed; Noordhof dots every possible "i" and crosses every possible "t". It contains an extended discussion of Humean supervenience, delves deeply into the theory of counterfactuals and the metaphysics of events which the theory presupposes, contains detailed discussion of so-called 'negative' causation, causal processes, the non-symmetry of causation, the relationship between causation and agency, causation and laws of nature, the metaphysics of chance, and much and more besides—more than I am able to concisely list in this review. Readers interested in any of these topics will find stimulating discussion in *A Variety of Causes*—though they may find the discussion somewhat daunting to consume. If I were to raise one concern with the book, it would be that it does little to assist readers who wish to 'skip ahead' to a certain topic without taking in the entirety of the vast theoretical apparatus developed over 500 pages, and without reading through detailed exegesis and criticism of alternative proposals from the literature. While the discussion is rich, not much of it is bite-sized.

In this review, I will focus on the theory of causation which forms the heart of the the book, and which is presented and defended in chapter 4. According to Noordhof, causation is not counterfactual dependence, but rather counterfactual dependence *modulo* a set of possible events, Σ. Some notation: let C and E be events which actually occurred, and let 'c' and 'e' be the propositions that C and E happen, respectively. For any set of events Σ, let σ be the proposition that some event in Σ happens—so that $\neg\sigma$ is the proposition that no event in Σ happens. Noordhof's full definition of Σ-dependence has bells and whistles to handle the nuances of probabilistic causation,[2] but if we focus on deterministic systems, this definition can be simplified. In that special case, we can say that

[2] Let '$\phi \ \Box\!\!\rightarrow_{x,t} \ \psi$' be true iff the mean objective chance of ψ just before time t in the most similar ϕ-worlds is x. And let t_E be the time at which E happened. Then, the full definition of Σ-dependence is this: E Σ-depends upon C ($C \notin \Sigma$) iff there are numbers h, l, with $h \gg l$, such that (i) $(\neg\sigma \wedge c) \ \Box\!\!\rightarrow_{h,t_E} \ e$, and (ii) for every time t, there's some $x \leqslant l$ such that $(\neg\sigma \wedge \neg c) \ \Box\!\!\rightarrow_{x,t} \ e$.

Σ-Dependence (Determinism) In deterministic systems, for any set of possible events Σ, E Σ-depends upon C ($C \notin \Sigma$) iff

(i) If C were to occur without any of the events in Σ occurring, then E would happen

$$(\neg\sigma \wedge c) \ \Box\!\!\rightarrow e$$

and

(ii) If neither C nor any of the events in Σ were to occur, then E would not happen.

$$(\neg\sigma \wedge \neg c) \ \Box\!\!\rightarrow \neg e$$

Causation is then defined as follows.

Causation C is a cause of E iff there is an appropriate set of possible events Σ such that:

(1) E Σ-depends upon C;

(2) there is no superset of Σ, Σ^*, such that both

(2a) E Σ^*-depends upon C, and

(2b) there is a non-actual event (an event which did not actually happen), $N \notin \Sigma^*$, such that E Σ^*-depends upon N.

If the set of events Σ can be used to show that C is a cause of E in *Causation*, say that Σ is a *witness* to C's being a cause of E.

I say that Σ must be an "appropriate" witness (the term is mine, not Noordhof's). What it takes for Σ to be appropriate is specified on page 139—though, I must confess that I am not sure I have understood the definition. But thankfully, I don't think we have to get distracted with the details. For Noordhof's goal in restricting the account to "appropriate" witnesses is to deal with situations in which the witnesses failing to occur would, *on its own*, change the times at which E may occur. So it appears that, so long as 'turning off' the events in Σ doesn't make any difference to when E might occur, Σ will be an appropriate witness.

We can sidestep issues having to do with indeterminism and issues related to the times at which the effect may happen by focusing on very simple deterministic systems in which each event can only happen at a certain time. So I'll limit my attention to systems of neurons, connected by stimulatory and inhibitory synapses, where every neuron is only able to fire at a very specific time. For instance, consider the case of *preemption* shown in figure 1. In figure 1, each circle represents a neuron, which can either fire or not fire at the time written underneath it. If a neuron fires at its designated time, then it is coloured grey. If it remains dormant, then it is coloured white. Arrows represent stimulatory synapses, whereas circular-headed connections (like the connection between C and B) represent inhibitory synapses. If C fires at t_1, then D will certainly fire at t_2, and B will certainly not fire at t_2—whether or not A fires. And E will fire at t_3 exactly if either B or D fires at t_2.

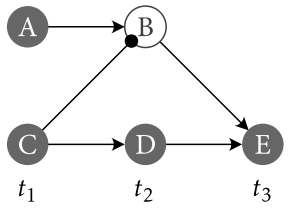

Figure 1: *Preemption*

(I am using non-italicised uppercase letters like 'N' to name neurons, and italicised uppercase letters like 'N' for the event of the neuron N firing. Lowercase italicised letters like 'n' therefore stand for the proposition that N fired.)

In *preemption*, C is a cause of E. However, A is a preempted *backup* cause of E—were it not for C, A would have been a cause of E. The presence of this preempted backup means that E does not counterfactually depend upon C. For, were C to not happen, A would have caused E to happen.

Nonetheless, Noordhof's counterfactual theory is able to tell us that C is a cause of E, if we use the witness $\Sigma = \{A\}$. For even though E does not depend upon C, it does $\{A\}$-depend upon C,

$$(\neg a \wedge c) \mathrel{\Box\!\!\rightarrow} e$$
$$\text{and} \quad (\neg a \wedge \neg c) \mathrel{\Box\!\!\rightarrow} \neg e$$

And, moreover, there's no superset of $\{A\}$, Σ^*, such that E Σ^*-depends upon

both C and some non-actual event $N \notin \Sigma^*$. So *Causation* rules that C is a cause of E.

To understand the reason for condition (2) in *Causation*, notice first that A *also* satisfies condition (1). For E $\{C\}$-depends upon A,

$$(\neg c \wedge a) \mathrel{\Box\!\!\rightarrow} e$$
$$\text{and} \quad (\neg c \wedge \neg a) \mathrel{\Box\!\!\rightarrow} \neg e$$

Condition (2) uses the fact that there is a 'gap' in the potential causal process leading from A to E to distinguish A from C. The 'gap' comes when B does not occur. Because this 'gap' exists, there is a non-actual event—namely, B—upon which E $\{C\}$-depends. That is:

$$(\neg c \wedge b) \mathrel{\Box\!\!\rightarrow} e$$
$$\text{and} \quad (\neg c \wedge \neg b) \mathrel{\Box\!\!\rightarrow} \neg e$$

So, even though condition (1) is satisfied, condition (2) is not. So $\{C\}$ is not a witness to A being a cause of E. In general, in order for Σ to be a witness to C being a cause of E, E must Σ-depend upon C, and must *not* Σ-depend upon any events which didn't actually occur.

But condition (2) of *Causation* doesn't just say that you can't have E Σ-depend upon a non-actual event. It also says that you can't have E Σ^*-depend upon a non-actual event, for any *superset* $\Sigma^* \supseteq \Sigma$. To appreciate why this additional strength is included, consider the system of neurons shown in figure 2. In this case, there are *two* 'gaps' in the backup processes leading from A to E. That is: neither B nor F happen. And because B and F would together symmetrically overdetermine E, were they to both happen, E does not $\{C\}$-depend upon either of these non-actual events individually.[3] But Noordhof notes that E *does* $\{C, F\}$-depend upon the non-actual event B, and E *does* $\{C, B\}$-depend upon the non-actual event F. So the additional strength of condition (2) allows Noordhof to deal with cases of preemption like this, as well.

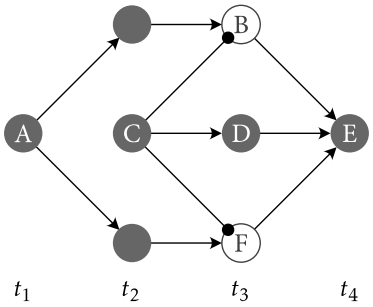

Figure 2: A is a backup which would have symmetrically overdetermined E, were it not for C.

[3] See Choi (2002) and Noordhof (2002).

It seems that the motivating idea behind Noordhof's theory is this: C can be a cause of E without E counterfactually depending upon C when there are other, would-be causes. (A would-be cause is just something that would be a cause, were C to not happen.) Would-be causes can sever the counterfactual relationship between E and C. But, if we 'turn off' these would-be causes of E, we should be able to restore counterfactual dependence between E and C without creating any new causal process leading from C to E. Condition (1) of *Causation* is meant to require that E depends upon C when any would-be causes of E are 'turned off', and condition (2) is meant to require that 'turning off' these would-be causes of E doesn't create any *new* causal process between C and E.

This approach to thinking about causation has much in common with the approach of authors like Hitchcock (2001), Woodward (2003), and Halpern & Pearl (2005). One important difference lies in the fact that, whereas these authors allow you to check for causation by checking for counterfactual dependence while 'holding fixed' whether certain other events occur or not, Noordhof only allows you to check for causation by

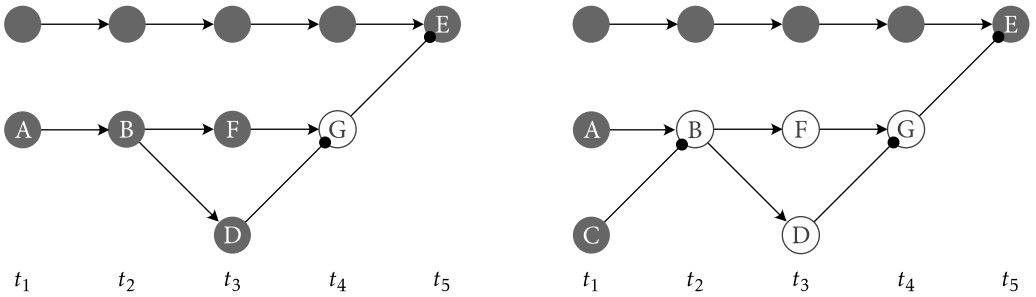

(a) A 'short circuit'

(b) A prevented 'short circuit'

Figure 3: Noordhof's theory says that, in figure 3a, neither A nor B is a cause of E. And, in figure 3b, it says that C is a cause of E.

checking for counterfactual dependence while 'holding fixed' the *non*-occurrence of events. In some cases, Noordhof's approach seems to do better. Consider, for instance the neuron system in figure 3a, which Hall (2007) calls a 'short circuit'. If we 'hold fixed' the occurrence of F, then whether E fires will counterfactually depend upon whether B does. So the theories of Hitchcock, Woodward, and Halpern & Pearl will say that both A and B are causes of E. Noordhof's theory disagrees, denying that either A or B is a cause of E. This strikes many of us as the right verdict (though some disagree).

However, consider what Noordhof says about the neuron C in figure 3b. This neuron system is exactly like the one in figure 3a, except that C prevents B from firing. In *this* case, E will counterfactually depend upon C when we hold fixed the *non*-occurrence of D. So Noordhof's theory will tell us that C is a cause of E.

More carefully, in figure 3b, E $\{D\}$-depends upon C,

$$(\neg d \wedge c) \boxdot\!\!\rightarrow e$$
$$\text{and} \quad (\neg d \wedge \neg c) \boxdot\!\!\rightarrow \neg e$$

Moreover, the only relevant non-actual events not in $\Sigma = \{D\}$ are $B, F,$ and G. But, had any of those events occurred without D, they would have *prevented* E from occurring. So E does not $\{D\}$-depend upon any of those non-actual events. Nor does it Σ^*-depend upon $B, F,$ or G, for any superset $\Sigma^* \supseteq \{D\}$.

Both intuitively and according to Noordhof's theory, B has no effect on E in figure 3a. But Noordhof's theory tells us that, in figure 3b, C is able to cause E by preventing B. This seems like the wrong verdict. If B has no effect on E when it happens, and the only thing C does is prevent B from happening, then C should not count as a cause of E.

Or consider the neuron system shown in figure 4. There, the connection between B and E is a *partially inhibitory* connection. If B fires, then E will need *two* stimulatory signals in order to fire. If, however, B doesn't fire, then E will only need one signal to fire. This neuron system feels similar to Hall's *short-circuit* from figure 3a. In figure 3a, A initiates a threat to E along one path by making F fire—which threatens to make G fire, which would prevent E from firing. But at the same time, A diffuses that very threat along another path by making D fire—thereby keeping G from preventing E from firing. Likewise, in figure 4, C initiates a threat to E along one path by

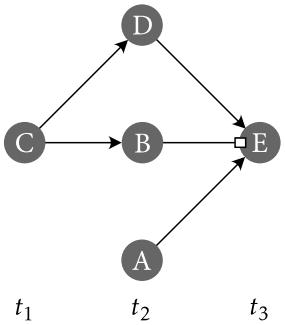

Figure 4: If B fires, E will need *two* stimulatory signals to fire. E $\{A, B\}$-depends upon C; so Noordhof's theory says that C is a cause of E.

making B fire—which threatens to keep E from firing. But at the same time, C diffuses that very threat along another path by making D fire—thereby keeping B from preventing E from firing. Just as I'm inclined to think that A is not a cause of E in figure 3a, I'm inclined to think that C is not a cause of E in figure 4.

Despite their similarities, Noordhof's theory distinguishes the two cases. In figure 3a, it says that A is not a cause of E. However, in figure 4, it says that C *is* a cause of E. For consider the witness $\Sigma = \{A, B\}$. If neither A nor B were to fire, then: (i) E would fire if C were to fire; and (ii) E wouldn't fire if C were to not fire.

$$(\neg a \wedge \neg b \wedge c) \mathrel{\Box\!\!\rightarrow} e$$
$$\text{and} \quad (\neg a \wedge \neg b \wedge \neg c) \mathrel{\Box\!\!\rightarrow} \neg e$$

Moreover, since there are no non-actual events to be considered in this neuron system (every neuron fires), there's no superset $\Sigma^* \supseteq \Sigma$ such that E Σ^*-depends upon a non-actual event.

The reader is left wondering whether, according to this theory, counterfactual dependence between distinct events is sufficient for causation. To appreciate why this isn't clear, consider the neuron system shown in figure 5. There, E $\varnothing$-depends upon C. Were C to fire, so too would E; and were C to not fire, neither would E. But the empty set is *not* a witness to C being a cause of E. For there is a superset of $\varnothing$, namely $\{A\}$, such that (2a) E $\{A\}$-depends upon C, and (2b) there is a non-actual event not in $\{A\}$, namely B, such that E $\{A\}$-depends upon B. Were neither A nor B to fire, neither would E; and, were A to not fire and B to fire, E would fire. Now, in this particular case, *Causation* does tell us that C is a cause of E. In this case, $\{B\}$ is a witness to C being a cause of E. But it was far from clear to this reader whether, in general, there will always be some set Σ which witnesses C being a cause of E when E $\varnothing$-depends upon C. (Matters are further complicated by the fact that, in general, events can happen at different times, requiring us to attend to the complexities of what it takes for a witness to be "appropriate", on page 139.)

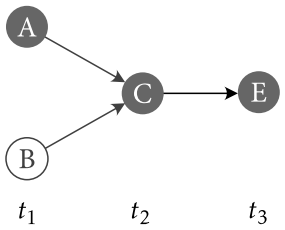

Figure 5: E $\varnothing$-depends upon C, but there is a superset of $\varnothing$, namely $\{A\}$, such that E $\{A\}$-depends upon the non-actual B.

In sum: I have a few concerns and lingering questions about the theory of causation defended in chapter 4. But I should note that many of the book's theses and arguments are largely independent of the finicky details of chapter 4. The broad defence of a counterfactual approach to causation which the book provides could be paired with any of a quite large number of counterfactual theories. Readers interested in counterfactual approaches to causation will find much to ponder over and learn from.

References

Choi, Sungho. 2002. "The 'actual events' clause in Noordhof's account of causation." In *Analysis*, **62** (1): 41–46.

Hall, Ned. 2007. "Structural Equations and Causation." In *Philosophical Studies*, **132** (1): 109–136.

Halpern, Joseph Y. & Pearl, Judea. 2005. "Causes and Explanations: A Structural-Model Approach. Part 1: Causes." In *The British Journal for the Philosophy of Science*, **56**: 843–887.

Hitchcock, Christopher. 2001. "The Intransitivity of Causation Revealed in Equations and Graphs." In *The Journal of Philosophy*, **98** (6): 273–299.

Noordhof, Paul. 2002. "Sungho Choi and the 'actual events' clause." In *Analysis*, **62** (1): 46–47.

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