Erschienen in: Philosophy of Science ; 86 (2019), 4. - S. 773-784
https://dx.doi.org/10.1086/704982

Reply to Jim Woodward's Comments on Wolfgang Spohn's *Laws of Belief*

Wolfgang Spohn*

This is one of a pair of discussion notes comparing some features of the account of causation in Wolfgang Spohn's *Laws of Belief* with the "interventionist" account in James Woodward's *Making Things Happen*. This note locates the core difference of the accounts in the fact that Woodward's account follows an epistemological order, while Spohn's follows a conceptual order. This unfolds in five further differences: (i) type- versus token-level causation, (ii) reference to time, (iii) actual/counterfactual intervention versus epistemic/suppositional wiggling, (iv) a circular versus a circle-free conception of the circumstances of a direct causal relation, and (v) absolute versus model-relative causation.

Woodward started writing on causation and explanation in 1979 (Woodward 1979), culminating in, and by far not ending with, his formidable book *Making Things Happen* (2003). I started writing on causation in 1978 (Spohn 1978, chap. 3), so far culminating in chapters 14 and 15, or 130 pages, in Spohn (2012). Our accounts of causation look similar; both seem to be variants of causal Bayes net theorizing (which was anticipated in Spohn (1978, 1980). Both look dissimilar (by Woodward's reliance on structural equations and my reliance on ranking theory), in ways that are of subordinate importance. And both are dissimilar in important ways that may be less obvious. Therefore, it may be useful to give an easily accessible description of those differences, which is perhaps an expedient companion to those many pages. Here, I am giving this description from my point of view, in the course of which I also respond to various comments by Woodward in his discussion note in this issue (2019).

The obvious difference is that Woodward states his account in terms of structural equations, or he starts illustrating it with them from the beginning in Woodward (2003, sec. 2.2). The equations generate or are represented by

*To contact the author, please write to: Department of Philosophy, University of Konstanz; e-mail: Wolfgang.Spohn@uni-konstanz.de.

a causal graph, and they are richly amended by statistical methodology, thus resulting in causal Bayes nets. In this way his account is well aligned with current scientific methodology. By contrast, I present my account in terms of my idiosyncratic ranking theory. This has a reason, of course, which is dear to me. I came to think of ranking theory in Spohn (1983), precisely with the aim of developing an account of deterministic causation—the kind of causation thinkers were thinking about for millennia—in perfect parallel to accounts of probabilistic causation, which at that time seemed superior and more sophisticated. These probabilistic accounts were more successful because they could use a more adequate notion of (conditional) independence. Ranking theory shares this success, while the structural equations approach still does not in my view.

The core point here is that the independence notion provided by structural equations is basically logical or functional dependence. Even with counterfactual dependence, one does not get anything like the graphoid axioms (cf. Spohn 1978, sec. 3.2; Pearl 1988, 82–87), which hold for conditional probabilistic (or ranking) independence. An important symptom of the relevant difficulties is the treatment of symmetric causal overdetermination, which is not satisfactorily dealt with by counterfactual theories and which has provoked at least four different versions of the ominous condition AC2(b) of the structural equations account of actual causation. The most recent one is in Halpern (2016, 22–26), which differs from his previous ones and also from AC*2 in Woodward (2003, 84). (For a more extensive comparison of the structural equations and the ranking-theoretic approach to causation, see Spohn [2010].)

This difference between the structural equations and the ranking-theoretic approach is neglected by Woodward's discussion (2019). Rightly so; it is, in a way, superficial. I have emphasized that ranking theory may be replaced throughout by probability theory and that all definitions and theorems should and can be preserved, thus resulting in my proposal for a probabilistic theory of causation. Likewise, structural equations can be probabilistically amended. And then the difference is largely gone, opening the view on the more basic differences.

There is also a salient difference in style. I recall when first reading Woodward's (2003) book that it almost felt like reading ordinary language philosophy. That is praise. I was raised with this kind of philosophy and love it. The book is a most considerate investigation into the conception of causation of laymen and scientists. By contrast, my chapters are largely a formal exercise in rigorous theorizing (and thus much less pleasant to read). Again, one might say that this is a superficial difference, while we are united in conceptual analysis.

However, this already moves me to what I take to be our core difference. Woodward's conceptual analysis is very close to scientific practice (this is

why it is so instructive for scientists of various sorts), while mine is very much *more geometrico*. Or in more substantial terms: I think that Woodward's investigation in effect rather follows an *epistemological order*. At least his discussion note (Woodward 2019) comments on my account of causation distinctly from an epistemological point of view. By contrast, I try to keep to a *purely conceptual order*, without regard to its epistemological consequences. Both orders are legitimate, important, and informative. So I have nothing to criticize about Woodward's order. But they entail quite a number of further differences. I would like to explain five such differences.

(1) A first conspicuous difference is that Woodward deals with *type-level* (I say, generic) causation while I treat *token-level* (singular) causation. The reason is simple. The natural and the social sciences are interested in laws or invariances, and if they have causal form, they concern type-level causation. Woodward addresses those sciences. (One should never forget, though, that there are large fields, law and history, e.g., that are primarily concerned with token-level causation.) By contrast, my entry to causal theorizing was causal decision theory (CDT), and this is always about single decision situations and hence about token-level causation. Woodward (2019) discusses this contrast in sections 2 and 3.

The point now is that Woodward is thereby oriented at the epistemological order. At least in the sciences we first try to find out about the causal regularities (even if only statistical ones) and then apply them to the single cases. This agrees with his strategy to explicate actual causation only in a second step, namely, relative to given structural equations or causal laws (see Woodward 2003, 74–86).

Woodward cites the example that it is statistically overwhelmingly confirmed that smoking is a contributing type-level cause of lung cancer, while it may be very difficult to confirm that Jones's smoking for 30 years caused his lung cancer. Well, that is how the tobacco industry argued for a long time: there is no conclusive evidence in the single case. I would think, if the type-level claim is well confirmed, the token-level claim about Jones is prima facie equally well confirmed. Still, the example supports Woodward's case.

However, this cannot be the conceptual order. What else could causal regularities be but generalizations (or perhaps averages) of singular causal claims? In my view, the strict conceptual order can only proceed to first understand the latter in order to understand the former. This is the order I attempt to pursue.

The point is emphasized by the fact that type-level causation is always about causal dependence between (type-level) variables, while token-level causation is about causation between events or singular facts. What is usually overlooked is that there also are token-level variables and causal dependencies between them. In any case, it seems clear to me that causal dependence between (token- or type-level) variables is causation between some

realizations of those variables and thus derivative on causation between facts or events.

This difference is one that I do not only have with Woodward. Since 1978 I have spoken about token-level causation and am hence at cross-purposes with the mainstream on probabilistic causation (e.g., Cartwright 1979). It took me a long time to realize how crucial this difference is. Cartwright acknowledged it only in Cartwright (1989, 9), where she also came to grant priority to token-level causation.

I confess that I have no more to offer to type-level theorists like Woodward than the crude suggestion that statements about causal relations between type-level variables are generalizations or averages of causal relations between the corresponding token-level variables. This may well be too crude. The type-level theorist has the reverse problem, though. It is quite unclear what his statements tell about the single case. The problem is analogous to the old issue about the relation between statistical probabilities and single-case propensities, which seems as dim as ever.

(2) This point directly entails the next crucial difference (also discussed by Woodward [2019, sec. 4]). Since I start with causation between singular facts, I can assume that those facts are temporally located and that causes always precede their effects. (There are difficulties with simultaneous causation, and I take backward causation to be a *contradictio in adjecto*.) Of course, this reference to time considerably facilitates theorizing. For instance, there is only one causal graph respecting the temporal relations between token-level variables and the conditional probabilistic dependencies between them, while the latter alone, without reference to time, are often compatible with many causal graphs.

The problem is, as Woodward rightly states, that on the type level of causal theorizing temporal relations are rarely specifiable. My parents' income precedes my educational status, yet there is no temporal and only a causal relation between parental income and educational status in general. Statistical data reveal a lot of conditional statistical dependencies between type-level variables, and the causal theorist can proceed only from them. However, that is the dictate of the epistemological, not of the conceptual, order.

Woodward is right in pointing to the extent to which causal theorizing gets along without reference to time. This is indeed remarkable and shows how far one can get in the epistemological order. From the purely conceptual point of view, however, I do not see why I should burden myself with such restrictions. Conceptually, the relation between time and causation has been fundamental throughout history (even though it is contested how exactly to state it).

(3) The story continues. The notion of *intervention* is at the center of Woodward's account of causation (see Woodward 2003, chap. 3; 2019, sec. 2). In the epistemological order this is perfectly justified. *Actual* intervention produces by far the best and most direct confirmation of causal relations. If

I wiggle the variable *X* at will, and the variable *Y* wiggles accordingly, then *Y* obviously causally depends on *X*. (The details are extensively elaborated by Woodward.) However, as he is well aware, actual intervention is rare. Therefore, he tries to generalize this epistemological virtue by turning to *counterfactual* interventions, which can be applied to each causal relationship. (Counterfactually, I can also shift the moon and see what happens to the tides.)

I perfectly agree with this move. However, I feel that the advocates of this move—who are many, not only Woodward—grossly underrate its size and risk. In my view, the very rich and confusing literature on conditionals and counterfactuals only allows the conclusion that their truth evaluability is very dubious or shallow. For instance, do we have a robust realistic notion of a similarity ordering between worlds as we have of tables or electrons? Surely we have the peremptory intuition that at least some counterfactuals are true, and Woodward (2003, 118–23) makes as strong a case as possible that the narrow, although not well-delineated class of interventionist (non-backtracking) counterfactuals belongs to them. Indeed, if we could start from true causal laws (or structural equations), the counterfactuals derived from them inherit their truth evaluability. However, in my conceptual order this start is illegitimate, and the doubt about their truth evaluability reversely spreads to causal laws.

So, if we want to have a general account of the conditional and counterfactual idiom, we must not presuppose its truth evaluability and then explain it in some shallow way or find excuses why it often does not work properly. Rather, we must try to find some other analytical starting point (e.g., expressivism) and then try to explain how some conditionals can emancipate from that starting point and acquire proper truth conditions. At least, that is the strategy I propose in Spohn (2015).

What does this entail for causation? My basic explication is that *A* is a *direct cause* of *B* if and only if (iff) *A* and *B* are singular facts, *A* precedes *B*, and *A* is positively relevant to *B* given the circumstances. (And then I go on to define causation as the transitive closure of direct causation and defend this step against recent criticism [see Spohn 2012, sec. 14.11–13]. I will not substantiate this issue here.) The crucial point is that the only workable notion of positive relevance that I can find is an epistemic one, either in terms of subjective probabilities or of ranks (or possibly of other degrees of belief). Thus, my notion of direct causation is thoroughly epistemically relativized. This is our Humean heritage, which we must accept and not reject in my view.

Thus, in my view, as we have no choice but to start from this epistemic relativization, I share the intention to go beyond it. Somehow, we have to turn what is relative to our minds into an objective feature of the world. In other words, I adhere to 'projectivism', which Woodward (2003, 118–23) finds confused and misguided. Indeed, the projectivist is always in danger of talking in perverted ways. This is due to the causal character of the projectivistic metaphor: without projector nothing is projected, so without the

human mind, there is no causation in the world. That is absurd and rightly attacked by Woodward. But it is a misunderstanding of projectivism, which does not make any counterfactual claims. How else, though, is the metaphor to be understood? That is an inveterate difficulty. I try to solve it in Spohn (2012, chap. 15) by what I describe as the 'objectivization' of the epistemically relativized causal relations, which is, as far as I see, not subject to Woodward's criticism of projectivism. For me, this is still the most important issue about causation and one on which Woodward and I are clearly at odds. However, it is too large to be further pursued here.

The only point I want to make is that once we move from actual to counterfactual wiggling, we have in effect moved to epistemic or suppositional wiggling, as embodied in my explication. Positive relevance of *A* to *B*, epistemically interpreted, says nothing but that variation of degree of belief in *A* covaries with variation of degree of belief in *B* (given the circumstances). Hence, the counterfactual theorist is in fact perilously close, in my view, to the Humean heritage. We have here another point of similarity between Woodward and me paired with disagreement underneath.

(4) This brings me to the next issue. Whatever the relevant kind of wiggling, we must keep the *circumstances* fixed. What does this mean? I argue that the circumstances of *A*'s being a direct cause of *B* consist of the entire history of *B* without *A*. Conceptually this is crucial, since it frees the above explication of conceptual circles; circumstances are thereby explained only with reference to time. But this is an epistemological disaster, as Woodward rightly remarks. It seems then that we can never affirm causal relations because of that reference to the entire history. This is another facet of my basic theme.

I can offer some consolation. Given my explication and some mild auxiliary assumptions, the circumstances of *A*'s being a direct cause of *B* can be reduced to consist of all the other direct causes of *B*. This was already suggested by Cartwright (1979), and she inferred, wrongly in my view, that we are thus entangled in an inextricable conceptual circle. The reducibility means that the rich (circle-free) and the reduced (circularly defined) circumstances provably yield the same causal relations. So, it suffices to keep fixed only those other causes. This is what we try to do in actual experimentation. Our painful experience, though, is that we were often wrong with our guesses as to those other causes. Conceptually, we are guaranteed to be on the safe side only with the rich circumstances. Epistemologically, however, we do, and have to, proceed without any such guarantees.[1]

1. In temporally ordered Bayes nets, the same holds. For temporally located token-level variables *X* and *Y*, we can define *X* to be a parent of *Y* iff *Y* probabilistically depends on *X* given all the other variables in the past of *Y*. And then we can prove the causal Markov condition (without auxiliary assumptions), i.e., that *Y* is independent from its past and indeed from all its nondescendants given all its parents.

(5) I can offer another consolation (which Woodward [2019] discusses in sec. 5). We will see, though, that it is a Greek gift. When I said that the circumstances consist of the entire past of the direct effect without the direct cause, I did not literally mean the entire past but only the entire past insofar as it is represented within the causal model or its set $\boldsymbol{V}$ of variables. This is something in our grip. Clearly, though, it makes my notion of direct causation model relative. Woodward (2003, 55) makes a similar move when saying that Y directly causally depends on X iff an intervention on X would make a difference for Y provided all other variables in $\boldsymbol{V}$ are kept fixed. (Since he cannot refer to a temporal order in $\boldsymbol{V}$, he must keep fixed all other variables in $\boldsymbol{V}$, not only those in the past of Y.) He thus accepts the model relativity of direct causation, that is, of the distinction between direct and indirect causation. (He does not do so, though, for causation simpliciter; see below.)

That is no surprise. Of course, the direct/indirect distinction is model relative. A larger model may spell out the steps mediating what appears to be a direct causal relationship within a smaller model. However, the model relativity implied by my explication runs deeper (that is the Greek gift). An apparent causal relation may turn out spurious in the larger model, or it may show up only in the larger and not yet in the smaller model (where the latter can occur only when faithfulness is violated). In other words, Simpson's paradox cannot raise its head within a given causal model because no variable in the model is left for further conditioning that could reverse the probabilistic dependencies. But by looking outside the model or enlarging it, we may fully run into the paradox again.

This raises Woodward's suspicion. It appears that I am not talking about causation, after all, but only about conditional dependence. (As far as I know, so-called Granger causality has fallen into disgrace in economics, precisely for the same reason; it is said to provide only a method of forecasting rather than of causal inference. As such it is still appreciated. I find this assessment unjustified, as I am going to explain.)

So, all the conditional dependencies or (partial) statistical correlations within the model cannot prove real causation. This is true. This entails that causation is a *model-transcendent* notion. This insight can hardly be overemphasized. Of course, causation can be represented in models. But it cannot be ascertained within them; it always refers to things not represented in them. So, obviously my model-relative notion of causation does not capture the model-transcendent or absolute notion, which all causal theorists including Woodward intend. In short, I am missing the topic.

This criticism is well taken. However, this model transcendence presents a substantial problem, which I would like to discuss in the rest of this note. How do others try to fulfill their intention? Let me look at three attempts.

The original causal Bayes net theory of Spirtes, Glymour, and Scheines (2000) may appear model relative as well. However, they do not really say

what causation is. In the terms of Glymour (2004), they prefer the 'Euclidean' method of specifying only axioms (causal Markov, minimality, and possibly faithfulness) that a causal model has to satisfy. Those axioms, though, hold only under an important assumption: that the model is *causally sufficient* in the sense that every (direct) common cause of two variables in the model is in the model as well (the exact definition is slightly more complicated; see Spirtes et al. 2000, 45). That is a very strong assumption, and it is clearly model transcendent. How ever to ascertain it? The Emperor of China may exert all kinds of hidden influences. Who knows? Maybe he is a common cause of two variables in the model? Maybe God is invariably so, as occasionalism originally intended? Of course, Spirtes et al. (2000, 44) are right when saying that we are justified in disregarding such weird hypotheses. Scientists have great but often failing skills in selecting reasonable models. Still, this assumption is the entry of the unavoidable model transcendence into their theory.

Spirtes et al. were aware of the strength of that assumption and were not happy with it. Substantial parts of their subsequent theorizing are occupied with how much of causal inference about hidden, latent, or unmeasured variables is still possible when the assumption is weakened. This is not the place to go into this. It is clear that these are most interesting ways to get a hold on the model transcendence, which, however, are bound to be only partial.

Another example is Woodward's notion of an intervention, which is model transcendent as well. For him, as mentioned, Y directly causally depends on X iff there is an intervention I on X that changes (the probabilities of) Y while keeping the other variables in the model set $\boldsymbol{V}$ fixed. But such an intervention variable I must itself satisfy various partially causal conditions. This circular characterization of causation and intervention has been criticized, for example, by Glymour (2004) but is defended by Woodward as something to be expected. This is not my issue. My point is that one of those conditions requires that the intervention variable I is statistically independent of any variable that causes Y along some causal path that does not go through X (this is condition I4 in Woodward 2003, 98). Here, "any variable" must be taken as quantifying also about variables outside $\boldsymbol{V}$. Again, we have a reference to the don't-know-what's outside the causal model. But, of course, this condition is necessary. If there were such a statistical correlation or dependence, the intervention on X would not show X's influence on Y.

If I intervene at will, how could there be such a correlation? This rhetorical question points to an action theoretic justification of Woodward's requirement. Interventions are actions that fall under the competence of decision theory, indeed CDT, as most philosophers think, including me. A basic tenet of CDT (see Spohn 1978, 109–10) is that acts are exogenous, cut off from any causal ancestry. The intervention variable I is really *do* I, as conceived in the *do* calculus of Pearl (2000, 85–89) and as emphasized by Woodward (2003, 47) as well. Since the only effect of the intervention I is on the variable X,

each correlation *I* may have with *Y* must therefore run through *X*. Thus, Woodward's requirement is automatically fulfilled in the picture of CDT.

However, CDT represents the internal picture of the agent or intervener. She must think that her actions are cut off and uncorrelated in this way. But the external observer may find hidden correlations. Otherwise we could leave it, for example, to the physician to assign patients to the test group and the control group at will. However, the wisdom of randomized controlled trials (RCTs) is precisely not to rely on the physician's will because of those possible hidden correlations. And since Woodward takes the observer's external view, he is right in insisting on that requirement.

This brings me to the third way of accommodating the model transcendence of causation, which is offered also by Woodward in response to my worries. The whole point of the experimental methodology of RCTs is precisely to average out all those model-transcendent factors and thus to deliver reliable average type-level causal hypotheses. Is this not good enough?

Well, RCTs require a lot of care. The randomization must be carried out properly, one must check for all selection variables potentially built in into the experiment, and so on. Good experiments attend to all these things. Given this, there is no doubt that RCTs belong to the most advanced epistemological strategies we have in the social sciences, epidemiology, and so on.

However, I do not think that they can fully control what I call the model transcendence of causation. RCTs can properly randomize over the given population. If well done, they optimize internal validity of the experiment. However, as is well known, this is no guarantee of external validity, of the generalizability of the experimental findings to other populations. Still worse, is the given population, even if we take it to be the entire present mankind, representative for the behavior of all the variables outside the model? It is not even clear what this representativity could mean when the joint distributions of the variables change all the time. In practice we have nothing better to go for. In theory, though, such an assumption of representativity would be a big and presumably not well-founded step. Again we find the opposition between the epistemological opportunities and the conceptual requirements.

Does the original counterfactual analysis of causation, as introduced by Lewis (1973) and refined later on, not tackle this model transcendence straightaway? It may seem so. This analysis does not mention models at all and immediately refers to possible (grand) worlds, which Lewis takes in the most comprehensive, not transcendible sense. However, the downside of this procedure is its startling theoretical poverty. Surprisingly, this poverty seems of no concern within the tradition of the counterfactual analysis. How many theorems do we find there? What a contrast to the rich development in Pearl (2000), Spirtes et al. (2000), and Woodward (2003). Woodward completely agrees with this criticism. The reason is clear: as soon as we want to get to details, if only in a theoretical way, we have to describe the possible worlds,

with predicate and individual terms or with variables and their realizations, and so on. And thereby we capture only small worlds, that is, models, and nothing more. This is an unavoidable trade-off.

My conclusion from all this is this: we grasp causation within models, we intend to grasp absolute causation, and we have no clear idea about the relation between the two things and only very partial ways of grasping all the don't-know-what's not included in our models. My further conclusion is that we should reverse the priorities. We should not stare at absolute causation, deplore the insufficiencies of model-relative causation, and try to directly repair them. We should rather be happy to start with model-relative causation and then try to thereby approach absolute causation. This is the reversed strategy I recommend.

So, I propose to start with the above model-relative explication of direct and indirect causation. This allows rigorous and instructive theorizing about the model-relative notion. For instance, as indicated above, the model-relative versions of the causal Markov and the minimality condition simply turn out to be provable. Thus, we catch up with Spirtes et al. (2000). As Woodward notes, this is not absolute causation. Can we somehow approach it from this starting point?

I think so. Not directly, though. However, we can study how model-relative causation in smaller and larger models relates. That is, it may be hard to say how model-relative causation in a small model unfolds in a larger model (although not everything goes). We can precisely study, however, how model-relative causation in a larger model appears in a smaller model. I have only theorems about causal dependence between variables and none about causation between facts. Roughly, one of my theorems (see Spohn 2012, 395–98) is that, given faithfulness, causal dependence of Y on X in a smaller model means that in the larger model Y causally depends on X or shares a common causal predecessor with X. (Faithfulness is not a necessary condition, but I do not know of any weaker sufficient condition.) In a way, this is not surprising. It is very similar to Reichenbach's common cause principle, which is intuitively very convincing (and turns out provable as well, taken in a model-relative way; see Spohn 2012, 384–86).

The idea now is that causation in the larger model is a substitute for model-transcendent causation in the smaller model. Of course, it is still model-relative causation. But we may conceive of absolute causation as causation relative to the universal model containing all variables whatsoever. Admittedly, this universal model is absolutely fictitious and ill defined, just like Lewis's grand worlds, which we may postulate but can describe only very partially. Still, we may take the way in which causation relates within smaller and larger models as paradigmatic or indicative of the relation between model-relative and absolute causation and hence state, for example, that model-relative causal dependence is either causal dependence or having a common

causal predecessor in the absolute sense (see assertion 14.42 in Spohn 2012, 397).

This move makes none of our substantial epistemological problems with (the intended) absolute causation vanish; it does not even mitigate them. We always move within limited models, we always try to extend them to larger but still limited models, and we always face the issue of what this might tell us about full reality or absolute causation. It is this issue that has been tackled in the work mentioned above. And my hunch is that it can be stated in an equivalent way within my proposal. (However, this should be checked. Maybe my hunch is too optimistic; this would throw doubt on my proposal. Or maybe my framework allows improving on that issue in some ways.)

Still, my proposal contains a determinate conceptual strategy: first to give a rigorous account of model-relative causation and then to approach absolute causation in those model-relative terms. This is clearer and more consistent than trying to account for absolute causation through amending model-inherent theorizing by inexplicable model-transcendent conditions. In other words: it reestablishes priority of the Socratic over the Euclidean method.

I said in the beginning that both the epistemological and the conceptual orders are legitimate and important. They are indeed. Implicitly, though, my discussion note was a pleading for the priority of the conceptual over the epistemological order. We should strictly pursue the conceptual order first and not compromise it with the epistemological order, as I interpret Woodward as doing. Only then we can know what our epistemological problems are and attempt to solve them.

REFERENCES

Cartwright, Nancy. 1979. "Causal Laws and Effective Strategies." *Noûs* 13:419–37.

———. 1989. *Nature's Capacities and Their Measurement*. Oxford: Clarendon.

Glymour, Clark. 2004. Review of *Making Things Happen*, by James Woodward. *British Journal for the Philosophy of Science* 55:779–90.

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

Lewis, David. 1973. "Causation." *Journal of Philosophy* 70:556–67.

Pearl, Judea. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufmann.

———. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Spirtes, Peter, Clark Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search*. 2nd ed. Cambridge, MA: MIT Press.

Spohn, Wolfgang. 1978. *Grundlagen der Entscheidungstheorie*. Kronberg Ts.: Scriptor. Available at https://www.philosophie.uni-konstanz.de/ag-spohn/personen/prof-dr-wolfgang-spohn/books-lecture-notes/.

———. 1980. "Stochastic Independence, Causal Independence, and Shieldability." *Journal of Philosophical Logic* 9:73–99.

———. 1983. "Eine Theorie der Kausalität." Habilitation thesis, University of Munich. Available at https://www.philosophie.uni-konstanz.de/ag-spohn/personen/prof-dr-wolfgang-spohn/books-lecture-notes/.

———. 2010. “The Structural Model and the Ranking Theoretic Approach to Causation: A Comparison.” In *Heuristics, Probability and Causality: A Tribute to Judea Pearl*, ed. Rina Dechter, Hector Geffner, and Joseph Y. Halpern, 493–508. San Mateo, CA: Kauffmann.
———. 2012. *The Laws of Belief: Ranking Theory and Its Philosophical Applications*. Oxford: Oxford University Press.
———. 2015. “Conditionals: A Unified Ranking-Theoretic Perspective.” *Philosophers’ Imprints* 15 (1): 1–30.
Woodward, James. 1979. “Scientific Explanation.” *British Journal for the Philosophy of Science* 30:41–67.
———. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
———. 2019. “On Wolfgang Spohn’s *Laws of Belief*.” *Philosophy of Science*, in this issue.