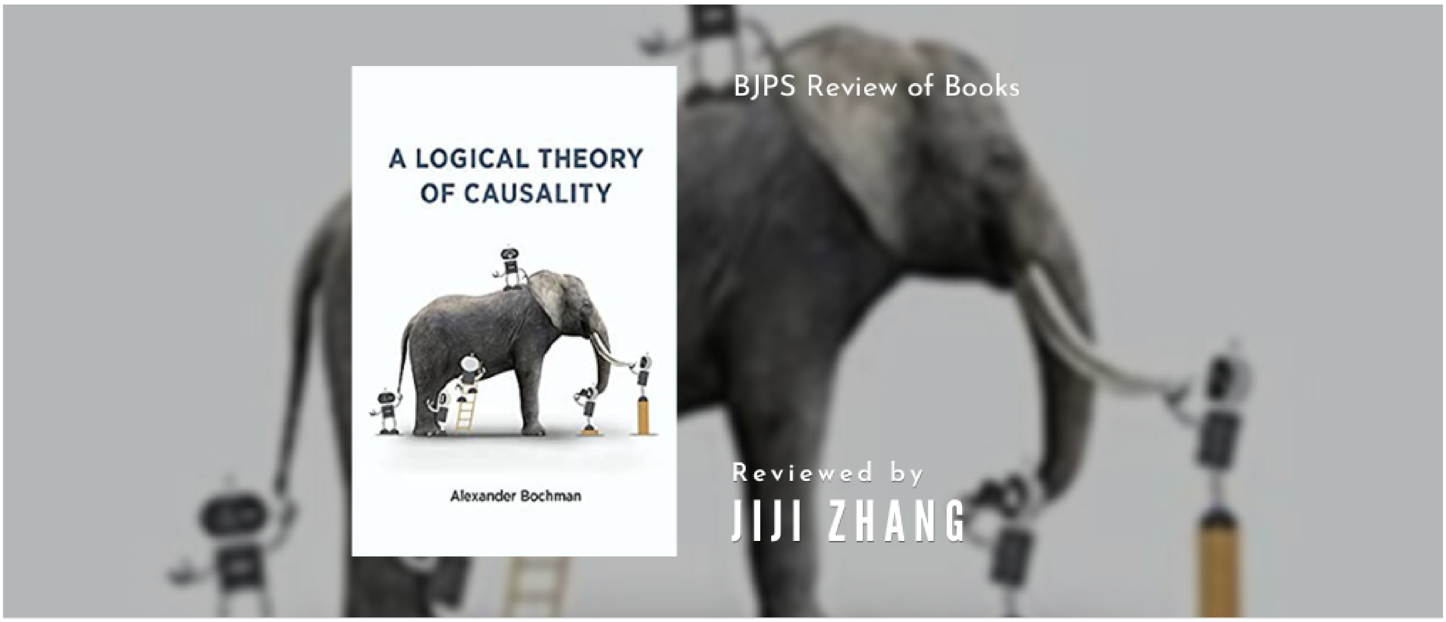

A LOGICAL THEORY OF CAUSALITY

Alexander Bochman

Reviewed by Jiji Zhang

A Logical Theory of Causality
Alexander Bochman
Cambridge, MA: MIT Press, 2021, £48.00
ISBN 9780262045322

This is a rich and illuminating book by a seasoned researcher in the field of artificial intelligence. Alexander Bochman is not only an expert in the logical approach to AI, especially knowledge representation and reasoning, but is also versed in the debates over causation in the philosophical and legal literature. Drawing on decades of vigorous work on the logical foundations of non-monotonic reasoning and their connections to causal reasoning, Bochman covers a broad territory in this book, ranging from classical logic and non-monotonic reasoning to systems of abduction and accounts of causation in a unifying non-monotonic formalism, and eventually to dynamic causal reasoning. Despite the dense formal results in many parts of the book, Bochman's writing is overall engaging and often insightful, with a pleasant dose of scholarly erudition.[1]

Another merit that facilitates reading is that this book is as modular as advertised in its preface. The four parts of this book, each consisting of three chapters, are cohesive to a good extent; the flow from any part to the subsequent one is smooth and well motivated. At the same time, the later parts are sufficiently self-contained that much therein can be studied fruitfully with no or limited need for the digestion of the earlier materials. For example, a reader who (like me) is most interested in the causal calculus and the logical approach to causation can safely skip Chapters 2 and 3 (I recommend reading Chapter 1 in any case), study Chapter 4 carefully (except for Sections 4.6 and 4.9) and Chapter 5 quickly (assuming some prior familiarity with structural causal models), and then jump to Chapters 8 and 9. Similarly, for someone who is primarily interested in dynamic causal reasoning, they can leap from Chapter 4 straight to Part 4.

Instead of giving a summary of the whole book, which I cannot do better than Bochman's own chapter-wise summary in the preface (pp. xiii–xvi), I will take advantage of this review to share some reactions to Bochman's accounts of general and actual causation in the framework of causal calculus, in which philosophers of causation are likely to take a special interest. What follows aims to give prospective readers a sense of Bochman's formalism, as well as the central ideas of his account, and to demonstrate the thought-provoking nature of these ideas.

The (propositional) causal calculus is built on classical Boolean logic. We begin with a standard language of the latter and add to it a conditional operator, $\Rightarrow$. The antecedent and consequent of the conditional are restricted to be Boolean formulas, so no nesting of $\Rightarrow$ will appear. Call a formula of the form $A \Rightarrow B$ a causal rule; a set of such rules forms a causal theory. Intuitively, a causal rule is intended to express a causal mechanism, as a structural equation in a structural causal model is intended to do (Pearl [2000]). Consequently, a causal theory can be used to describe or represent a causal system or situation as a structural causal model is used to do. Formally, Bochman shows in Chapter 5 that at least for Boolean variables the full machinery of structural causal models, including the modelling of interventions by sub-models, can be represented in his framework. More interestingly, Bochman argues in Chapter 8 that the additional expressive power of his framework carries a distinctive advantage in theorizing about actual causality. We will return to this issue below.

Below are some candidate derivation rules for the causal calculus, in which $\vDash$ denotes classical entailment:

(1) [Strengthening] If $A \vDash B$ and $B \Rightarrow C$, then $A \Rightarrow C$;

(2) [Weakening] If $A \Rightarrow B$ and $B \vDash C$, then $A \Rightarrow C$;

(3) [And] If $A \Rightarrow B$ and $A \Rightarrow C$, then $A \Rightarrow B \wedge C$;

(4) [Truth] $\mathbf{t} \Rightarrow \mathbf{t}$;

(5) [Falsity] $\mathbf{f} \Rightarrow \mathbf{f}$;

(6) [Cut] If $A \Rightarrow B$ and $A \wedge B \Rightarrow C$, then $A \Rightarrow C$;

(7) [Or] If $A \Rightarrow C$ and $B \Rightarrow C$, then $A \vee B \Rightarrow C$.

Bochman calls the relation expressed by $\Rightarrow$ (which he prefers to regard as an inference relation), a production inference relation if it satisfies 1–5, a regular one if it satisfies 1–6, a basic one if it satisfies 1–5 + 7, and finally, a causal one if it is both regular and basic. Accordingly, we can define the production, regular, basic, and causal closure of a causal theory, respectively, and the corresponding notions of equivalence between causal theories. For example, two causal theories are basically equivalent if they have the same basic closure.[2] Chapters 4 and 5 contain a wealth of interesting results to characterize these equivalences, which I suspect will turn out to be useful for causal discovery based on this formalism, a topic that is not addressed in this book.

For Bochman, general causation is to be captured by a causal inference relation. His account of actual or singular causation, on the other hand, follows an INUS or NESS (necessary element of a sufficient set) approach. It is defined relative to a literal causal theory, in which all causal rules are of the form that the consequent is a literal and the antecedent is a conjunction of literals. The central idea is this: A literal l is an actual cause of a literal m (relative to a literal causal theory and a world) just in case in the regular closure of the causal theory (more on this below), there is a causal rule $l \wedge L \Rightarrow m$ instantiated in the world (where L is a possibly empty conjunction of literals or **t**) but there is no $L \Rightarrow m$ as a causal rule. That is, l is a necessary or non-redundant part of an (instantiated) causally sufficient condition for m.

A contention stressed by Bochman is that general causation and actual causation have conspicuously different logical profiles: 'general causation [is] a purely *logical* notion that is described by an appropriate formalism of causal inference. In contrast, actual causation is already an explicitly *nonmonotonic* notion since it depends on the absence (non-provability) of certain causal rules' (p. 196). This is an interesting claim, but Bochman's apparent identification of general causation with his causal inference relation seems to be based on a crude understanding of the former. Bochman intends $A \Rightarrow B$ to express a general or type-level causal statement that A causes B (p. 79), but the postulates of his causal inference relation, especially [Strengthening] and [Weakening], suggest that it is more accurate to read $A \Rightarrow B$ as saying that A is causally sufficient for B.[3] I do not deny that a stipulative identification of general causation with causal sufficiency can be useful for some purposes, but in both scientific and everyday discourse, it is more typical to understand a statement of general causation as one of causal relevance. We say smoking causes cardiovascular diseases as a general fact even though smoking is by itself not causally sufficient, and we are reluctant to say smoking with glasses on causes cardiovascular diseases because we believe wearing glasses is not causally relevant. In fact, in either the standard INUS-style regularity account (for example, Baumgartner [2013]) or the 'remixed' account of Strevens ([2007]) that seems to have inspired Bochman, the concept of general causation is likely to also contain some notions of non-redundancy and be as explicitly non-monotonic as that of singular causation.

Another logical difference indicated by Bochman is that, as highlighted previously, actual causation in his account has regular inference as the underlying logic (that is, without the derivation rule [Or] but with the rule [Cut]), whereas general causation presumably follows causal inference. Again, this claim is at least misleading due to the conflation of general causation and causal sufficiency. More importantly, Bochman's choice of regular inference for actual causation is questionable, both because his reasons for resisting [Or] are less than compelling and because [Cut] is problematic.

Due to space limitation, I will consider here only one of Bochman's reasons for rejecting [Or]. The reason is illustrated by an example borrowed from Tim Maudlin (p. 190). Simplifying the example, suppose we have two causal rules: $p \wedge q \Rightarrow r$ and $\neg p \wedge q \Rightarrow r$, which represent two different mechanisms for producing r. Suppose the actual world is $\{p, q, r\}$, in which the first rule is instantiated. Is p an actual cause of r? Intuitively, Bochman suggests, it is. And since the rule $q \Rightarrow r$ is not derivable by regular inference, his account concurs. By contrast, if we allow the application of [Or], we get $(p \wedge q) \vee (\neg p \wedge q) \Rightarrow r$, and consequently also $q \Rightarrow r$, which would render p redundant for r. Note that the typical accounts based on structural causal Baumgartner (as well as representative regularity accounts such as Baumgartner [2013]) are in effect committed to the equivalence between $\{(p \wedge q) \vee (\neg p \wedge q) \Rightarrow r\}$ and $\{p \wedge q \Rightarrow r, \neg p \wedge q \Rightarrow r\}$. The option to distinguish them in Bochman's framework is therefore regarded as a notable advantage conferred by its extra expressive power.

Interesting as it is, this putative counterexample to [Or] seems to ride on a dubious intuition. If p as opposed to $\neg p$ cannot make any difference to r under any contingency, as is the case in the envisaged system, it seems perfectly reasonable to deny p as a (general or singular) cause of r. It is not the place to argue for this difference-making requirement in detail, but consider this question: What if $l \wedge L \Rightarrow m$ and $L \Rightarrow m$ represent two mechanisms, respectively? Why must we think that when there is a mechanism for L alone to produce m, there cannot be another, over-determining mechanism for l and L together to produce m? This consideration suggests that the whole INUS or NESS approach presupposes some difference-making condition on causation (or even on individuation of mechanisms).

Turn now to Bochman's acceptance of [Cut]. My challenge is related to a nice result Bochman proves in Chapter 5: basic equivalence ensures equivalence regarding intervention effects (Theorem 5.5), but regular or causal equivalence does not. In other words, in addition to the postulates for a production inference relation, applying [Or] will not conflate causal theories that can be distinguished by interventions, but applying [Cut] may. The simplest example to illustrate the latter is Example 5.7 (p. 132). The two causal theories, $\{p \Rightarrow q, p \Rightarrow r\}$ and $\{p \Rightarrow q, p \wedge q \Rightarrow r\}$, are regularly equivalent because [Or] allows the derivation of $p \Rightarrow r$ from the latter and [Strengthening] allows the derivation of $p \wedge q \Rightarrow r$ from the former. But they can be distinguished by setting q to be false by an intervention.

Exactly such an example was once used to challenge the non-redundancy condition in INUS (Zhang [2017]). Two professors decide whether a student's thesis passes or not according to the rule: a thesis passes (r) just in case P_1's verdict is positive (p) and P_2's verdict is positive (q). As it happens, the situation is such that P_2 is deferent to P_1 and always manages to copy P_1's verdict. It is thus described by the causal theory $\{p \Rightarrow q, p \wedge q \Rightarrow r\}$. In the circumstance where both verdicts are positive and the thesis passes, is P_2's verdict being positive a cause of the passing of a thesis? I think it is, with high confidence and unambiguous support from the structural–interventional accounts. Bochman's account delivers the wrong answer due to the derivation of $p \Rightarrow$

r by [Cut]. Interestingly, unlike the standard regularity accounts based on classical logic, Bochman can easily avoid this challenge by giving up [Cut]. This is, I think, an illustration of the advantageous flexibility of the framework.

I am therefore inclined to think that in Bochman's framework, the more adequate logic for actual causation (and for general causation as causal relevance) is after all basic inference, the inference relation that does not collapse interventional differences.

These critical remarks are not so much registering objections as illustrating the constructive thoughts a part of this book has provoked in the present reader. Much else is worth engaging with, especially the non-monotonic semantics for the causal calculus induced by a requirement of universal causation or explanatory closure. This book deserves to be studied by anyone who is interested in inferential approaches to causality and logical foundations of causal reasoning.

Jiji Zhang
Hong Kong Baptist University
zhangjiji@hkbu.edu.hk

Notes

[1] There are numerous useful references, but one mis-reference can be seriously misleading. The reference to (Halpern [2000]) on p. 128 should instead be to (Halpern [2013]), which is missing from the bibliography.

[2] I use this example because, unlike regular equivalence and causal equivalence, this notion does not seem to be explicitly defined in the book despite it playing an essential role in the results on intervention-equivalence in Chapter 5, Section 5.4.

[3] In the article introducing the language of causal theories, McCain and Turner ([1997]) gave a more modest reading of $A \Rightarrow B$: if *A*, then *B* is caused (not necessarily by *A*).

References

Baumgartner, M. [2013]: 'A Regularity Theoretic Approach to Actual Causation', *Erkenntnis*, **78**, pp. 85–109.

Halpern, J. Y. [2013]: 'From Causal Models to Counterfactual Structures', *The Review of Symbolic Logic*, **6**, pp. 305–22.

McCain, N. and Turner, H. [1997]: 'Causal Theories of Action and Change', in AAAI, *Proceedings of the 14th National Conference on Artificial Intelligence*, Menlo Park, CA: AAAI Press, pp. 460–65.

Pearl, J. [2000]: *Causality: Models, Reasoning, and Inference*, Cambridge: Cambridge University Press.

Strevens, M. [2007]: 'Mackie Remixed', in J. K. Campbell, M. O'Rourke, and H. S. Silverstein (*eds*), *Causation and Explanation*, Cambridge, MA: MIT Press, pp. 93–118.

Zhang, J. [2017]: 'On the Minimization Principle in the Boolean Approach to Causal Discovery', in C. Yang, K. Lee and H. Ono (*eds*), *Philosophical Logic: Current Trends in Asia*, Singapore: Springer, pp. 79–94.