Bayesian Nets and Causality: Philosophical and Computational Foundations, by Jon Williamson. Oxford: Oxford University Press, 2004. Pp. x + 250. H/b £74.50.

There is no doubt about the close connection between causality and probability. For instance, suppose that I throw a stone at a bottle, the stone hits the bottle, and the bottle shatters. In this case, my throw is a cause of the shattering of the bottle. Also it raises the probability of the shattering of the bottle, namely, P(the shattering of the bottle/my throw) > P(the shattering of the bottle). This observation motivates a long and strong tradition in the philosophy of causality that attempts to cash out causality in terms of probabilistic relationships. However, it turns out that the connection between causality and probability is much trickier than previously thought. There is a large degree of consensus among philosophers that a cause does not directly raise the probability of its effect nor is it connected to its effect via a chain of stepwise probability-raising relations.

Bayesian Nets and Causality is a product of Jon Williamson's ambitious project of providing a unifying picture of causality and probability. Very simply, with regard to probability, he supports objective Bayesianism that probabilities are an agent's rational degrees of belief fixed by her background knowledge. Meanwhile, with regard to causality, he advocates the epistemic view of causality that causality is basically a mental matter in the sense that a causal structure is part of an agent's mental representation of the world. Williamson takes it that an agent's causal beliefs are determined by causal and probabilistic constraints imposed by her background knowledge almost in the same way as an agent's rational degrees of belief are determined by causal and probabilistic constraints imposed by her background knowledge.

Bayesian Nets and Causality is a very well-written and organized book. Also it addresses a number of important metaphysical and computational issues concerning Bayesianism, the nature of causality, causal reasoning, and so on. No doubt it will be recognized as a very important contribution to the philosophy of probability and causality by a young distinguished philosopher.

Roughly speaking, the first half of *Bayesian Nets and Causality* is devoted to the philosophy of probability and the second half to the philosophy of causality. In the first half of the book, Williamson develops and defends the objective Bayesian interpretation of probability. Objective Bayesianism is the view that probabilities are an agent's rational degrees of belief, where an agent's degree of belief is rational iff (1) it does not incur a Dutch book; (2) it satisfies constraints due to our empirical knowledge; (3) it satisfies constraints due to our lack of further knowledge. On this view, given an agent's probabilistic and

causal background knowledge, only one degree of belief comes out rational. This means that when an agent's probabilistic and causal background knowledge is fixed, there is a unique degree of belief that she has to adopt, in which sense, on objective Bayesianism, probability is objective. By way of arguing for objective Bayesianism, Williamson touches on plenty of computational issues such as how our empirical background knowledge constrains our degrees of beliefs and how to exactly formulate the underlying idea of the principle of indifference.

One of the important issues Williamson addresses in this regard is the Causal Markov Condition that states that 'each variable is probabilistically independent of its non-effects, conditional on its direct causes' (p. 50). He argues that, when we interpret the Causal Markov Condition as concerning physical probability and causality, well-known counterexamples to the Principle of Common Cause such as Elliot Sober's and Frank Arntzenius's equally serve as counterexamples to the Causal Markov Condition because the latter implies the first (p. 52–7). On this ground and others, Williamson asserts that the Causal Markov Condition fails for a variety of interpretations of probability and causality. But it does not immediately follow from this that it is wrong altogether. Indeed, Williamson suggests that the causal interpretation of Bayesian networks that assumes the Causal Markov Condition can be justified by invoking his technique of translating an agent's probabilistic and causal background knowledge into precise quantitative constraints on her belief function and then determining the belief function from those constraints by maximizing entropy. And he goes on to show that the resulting belief function can be represented by a causally interpreted Bayesian net. From this he concludes that, on the objective Bayesian interpretation of probability, we can justify the causal interpretation of Bayesian networks, and therefore, the Causal Markov Condition.

Meanwhile, in the second half of the book that I think far more contentious than the first half, Williamson addresses many issues on causality. He first criticizes a variety of existing metaphysical views of causality such as the mechanism and probabilistic theories of causality. With regard to how to discover causal relationships, Williamson levels criticisms against Popper's hypothetico-deductive account and the recent inductivist account although later he suggests a modified hypothetico-deductive account the first step of which is accounted for by inductivism. Williamson's main proposal is his epistemic view of causality that he thinks is coherent with the objective Bayesian interpretation of probability. On this view, causality is mental in the sense that it is a feature of an agent's epistemic state—there are no such things as physical

causal relations—, yet it is objective in the sense that given that two agents have exactly the same background knowledge, they must have the same causal beliefs. Finally, Williamson attempts to extend his view to cover recursive causality.

Williamson backs up the epistemic view of causality by championing the claim that it is convenient to represent the world in terms of cause and effect and this convenience explains why we think in terms of cause and effect. Then in what sense is causality convenient? Williamson suggests 'the concept of cause is useful because a causal connection is (i) a reliable (though not fully reliable) indicator of a probabilistic dependence, … (ii) helpful for making strategic decisions' (p. 137). This suggestion is supported by the following two principles:

> **Qualified Causal Dependence:** Normally causal relations are accompanied by probabilistic dependencies
>
> **Strategy:** Normally, instigating causes is a good way to achieve their effects. On the other hand instigating effects is not normally a good way to bring about their causes.

On Williamson's view, these two principles fall short of constituting a probabilistic or agency account of causality because they are qualified by 'normally'. However, they provide a good enough explanation for why the concept of causality is useful. Moreover, for Williamson, they ensure that the causal beliefs that an agent ought to adopt are determined by her background knowledge because they lead to what Williamson calls 'Probabilistic to Causal Transfer' (p. 140).

A number of questions arise as to Williamson's epistemic view of causality, though. For a start, Williamson does not make clear what normality consists in nor does he offer any supporting arguments for his claim that although causality cannot be analysed in terms of probabilistic dependence or strategic effectiveness, still normally causality is accompanied by probabilistic dependence and strategic effectiveness. It may be plausible to suggest that there is a sense of normality in which Qualified Causal Dependence and Strategy hold for the actual world. But I am sceptical that there is a sense of normality in which they hold for all possible worlds. Among other things, it is not difficult to imagine possible worlds where redundant causality is commonplace, and therefore, normally causality does not come together with probabilistic dependence nor strategic effectiveness. For instance, imagine a possible world where an omnipotent God serves as backup potential causes of all events. In this possible world, for every event *e*, even if one of its causes were not to occur, God would come along and bring about *e*. Then it would not be the case that normally

causality is accompanied by probabilistic dependence and strategic effectiveness. If so, neither of Qualified Causal Dependence and Strategy is metaphysically true.

To be sure, Williamson is likely to reply that, in the possible world just described, causality would not exist on the ground that, given that the convenience of causality is spelt out by means of Qualified Causal Dependence and Strategy, it is not convenient to represent it in terms of cause and effect. But, in that world, my throwing a stone would bring about the shattering of the bottle in exactly the same way as my throwing a stone actually does. Therefore, I think, Williamson's possible response is counter-intuitive. To say the least, we can hardly be persuaded that causality would not exist in that possible world unless we are given strong supporting arguments for it. What if it turns out that neither of Qualified Causal Dependence and Strategy is metaphysically true? It is not clear to me how Williamson's view would fare in this case. To make it clear, I think, we need to delve into the metaphysical status of Williamson's epistemic view of causality.

Secondly, despite Williamson's effort, it is not clear at all how causal representations that, on the epistemic view of causality, constitute causality follow from our background knowledge. He holds that they are determined by an agent's causal and probabilistic constraints imposed by her background knowledge. But, given that Williamson denies the existence of physical causality, what constitutes causal knowledge that is supposed to supply causal constraints? Williamson's answer is that our knowledge of mechanisms, laws and temporal relations supplies causal constraints (p. 146). One minor point I want to make here is that it is not clear at all what Williamson means by 'mechanisms'. I think that mechanisms can be best understood to be causal mechanisms. Therefore, we cannot tell if there is a mechanism between two variables without relying on the concept of causality. Then Williamson is forced to say that given that there are no physical causal relations in the world, our knowledge of mechanisms cannot serve as a source of causal constraints.

More importantly, it remains to be clarified how our background knowledge imposes positive causal constraints. As Williamson says, information about mechanism, laws and temporal relations imposes only negative constraints. Then where do positive causal constraints come from? Here it may be suggested that positive causal constraints come from probabilistic constraints via what Williamson calls Probabilistic to Causal Transfer. Roughly speaking, it says that 'if background knowledge induces a strategic dependence from A to B, … then an agent should attribute the dependence to a causal relation A → B' (p. 140), where A strategically depends on B if A and B are

probabilistically dependent when intervening on A and controlling for B's other causes. But at least as far as the formation of an initial causal belief graph is concerned, we cannot draw causal constraints from probabilistic constraints. In the present context, the agent is supposed to be given no prior causal belief graph but form a causal belief graph from her background knowledge. However, the concept of strategic dependence is defined in terms of causality. To be specific, an agent cannot tell if A strategically depends on B unless she can identify what are B's other causes. This means that given that an agent is supposed to have no prior causal belief graph, she cannot determine if A strategically depends on B. But Williamson's Probabilistic to Causal Transfer is intended to transfer strategic dependencies to causal constraints. Therefore, if an agent cannot tell what strategically depends on what, no probabilistic constraints can be transferred to causal constraints via Probabilistic to Causal Transfer. Hence, at least when it comes to the formation of an initial causal belief graph, Williamson's Probabilistic to Causal Transfer fails to transfer probabilistic constraints to positive causal constraints. In short, the question of how our background knowledge imposes positive causal constraints still remains to be answered.

Finally, Williamson proposes that we can deny the existence of physical causality because we can explain the convenience of the concept of causality in terms of probabilistic dependence and strategic effectiveness. But, as Williamson says, this reasoning requires one additional assumption that 'Humans think in terms of cause and effect because of this convenience, not because there is something physical corresponding to cause which humans experience' (p. 135). Unfortunately, however, Williamson, in his book, gives no reasons to accept this crucial assumption. Of course, he criticizes a number of current accounts of physical causality in chapter seven. This, however, falls short of persuading us to accept that there is no physical causality at all. In addition, *prima facie*, it sounds quite plausible to say that the reason why it is convenient to represent the world in terms of cause and effect is that there is physical causality in the external world and that because of this we think in terms of cause and effect. As a result, I think that Williamson's epistemic view of causality is still not fully satisfactory.

Please see Jon Williamson's web page (http://personal.lse.ac.uk/willia11/) for responses to the questions raised in this review.

Programs in History and Philosophy of Science SUNGHO CHOI
College of Natural Sciences
Seoul National University

San 56-1, Sillim-dong, Kwanak-gu
Seoul, 151-742, Korea