



Bread prices and sea levels: why probabilistic causal models need to be monotonic

Vera Hoffmann-Kolss¹ 

Accepted: 26 January 2024
© The Author(s) 2024

Abstract

A key challenge for probabilistic causal models is to distinguish non-causal probabilistic dependencies from true causal relations. To accomplish this task, causal models are usually required to satisfy several constraints. Two prominent constraints are the causal Markov condition and the faithfulness condition. However, other constraints are also needed. One of these additional constraints is the causal sufficiency condition, which states that models must not omit any direct common causes of the variables they contain. In this paper, I argue that the causal sufficiency condition is problematic: (1) it is incompatible with the requirement that the variables in a model must not stand in non-causal necessary dependence relations, such as mathematical or conceptual relations, or relations described in terms of supervenience or grounding, (2) it presupposes more causal knowledge as primitive than is actually needed to create adequate causal models, and (3) if models are only required to be causally sufficient, they cannot deal with cases where variables are probabilistically related by accident, such as Sober's example of the relationship between bread prices in England and the sea level in Venice. I show that these problems can be avoided if causal models are required to be monotonic in the following sense: the causal relations occurring in a model *M* would not disappear if further variables were added to *M*. I give a definition of this monotonicity condition and conclude that causal models should be required to be monotonic rather than causally sufficient.

Keywords Causation · Causal Bayes nets · Monotonicity · Causal sufficiency · Variable choice

✉ Vera Hoffmann-Kolss
verahoffmann@gmx.de

¹ University of Bern Institute of Philosophy, Laenggassstrasse 49a, Bern 3012, Switzerland

1 Introduction

Contact with a person who has a SARS-CoV-2 infection increases the probability of acquiring such an infection oneself. Here, probabilistic dependence reflects an underlying causal dependence relation: transmission of SARS-CoV-2 pathogens is a causal process. However, it is well known that not all probabilistic or statistical dependence relations are indicative of causal relations. In some cases, probabilistic dependence relations occur because there are hidden common causes. Having a dry cough increases the probability of having an elevated temperature. But neither does the cough itself cause the fever, nor does the fever cause the cough. Dry cough and elevated temperature are symptoms of an infection, which is a common cause of both.

In other cases, non-causal probabilistic dependence relations cannot even be explained by latent common causes. There is a positive statistical correlation between my age and the average global temperature, but this correlation is not indicative of a hidden common-cause structure. One of the most famous cases of this type is Sober's example of the relationship between bread prices in England and the sea level in Venice. Both have increased steadily over the last two centuries, but there is no common causal process to explain this correlation (Sober, 2001, p. 332).

Probabilistic causal models rely on the observation that many causal relations lead to probabilistic dependencies. These dependencies can be described by a probabilistic model, which consists of a set of variables and a set of directed edges connecting these variables. The variables describe the (putative) causes and effects represented by the model. Given that not all probabilistic dependencies represent causal relations, a central challenge for this approach is to distinguish true causal relations from purely probabilistic dependence relations.

Therefore, causal models are required to satisfy several constraints. The two most prominent requirements are the causal Markov condition and the faithfulness condition (Spirtes et al., 2000, p. 29–31). However, other constraints are also important. Another condition that is crucial for distinguishing between probabilistic dependencies that indicate causal relevance relations and dependencies that occur for other reasons is the causal sufficiency condition. The basic idea of this condition is that causal structures should not overlook hidden common causes: if V is a set of variables constituting a causal model, then there must be no variable not included in V that is a direct cause of two or more variables in V (Spirtes & Scheines, 2004, p. 836; Spirtes, 2010, p. 1651; Zhang & Spirtes, 2011, p. 337). Accordingly, if the set of variables constituting a model includes the variables X and Y , then it should also include all direct common causes of X and Y . The causal sufficiency condition is usually regarded as a precondition for the application of the causal Markov condition, that is, the causal Markov condition is only applied to causally sufficient models (Spirtes & Scheines, 2004, p. 836; Spirtes, 2010, p. 1651; Zhang & Spirtes, 2011, p. 337).

In this paper, I argue that the causal sufficiency condition should be replaced by a monotonicity condition. Monotonicity in this context means that the probabilistic dependence relations in a causal model would not disappear if more variables were added. I argue that there are three ways in which the causal sufficiency condition is problematic: (1) A model can be adequate only if the variables it contains do not stand in non-causal necessary dependence relations, such as mathematical or conceptual

relations, or relations described in terms of supervenience or grounding. However, the causal sufficiency condition implies that variables standing in such relations must be included in the same causal model. (2) The causal sufficiency condition, as usually formulated, leads to an infinite regress. And if this regress is to be avoided, the condition must presuppose more causal knowledge as primitive than is actually needed to create adequate causal models. (3) The causal sufficiency condition (in combination with the causal Markov condition and the faithfulness condition) is too weak to account for causal structures that involve accidental probabilistic relations, such as the relation between my age and the average global temperature, or the relation between bread prices in England and the sea level in Venice. I argue that if causal models are required to be monotonic rather than being causally sufficient, they can handle all three problems better.

The paper is organized as follows: I begin with an overview of the causal modeling framework and the causal sufficiency condition (Sect. 2). I discuss the causal sufficiency condition in detail and show that it leads to the three problems mentioned above (Sect. 3). I then define the monotonicity condition and argue that it is superior to the causal sufficiency condition (Sect. 4). Finally, I compare my approach to Woodward's interventionist solution to the problem of accidental probabilistic relations (Sect. 5).

2 Probabilistic causal models and the causal sufficiency condition

A probabilistic causal model describes cause-effect relations as dependence relations between variables whose values can represent various things, such as 'the occurrence or non-occurrence of an event, a range of incompatible events, a property of an individual or of a population of individuals, or a quantitative value' (Hitchcock, 2023). Causal structures are represented by directed acyclic causal graphs consisting of two elements: (a) a set V of variables and (b) a set of directed edges connecting those variables. A sequence of variables $\{X_1, \dots, X_n\}$ is called a 'directed path from X_1 to X_n ' iff for any i with $1 \leq i < n$, there is a directed edge from X_i to X_{i+1} and all edges point in the same direction. If Y is a variable in V , such that there is a directed edge from Y to X , then Y is called a 'parent of X '. If there is a variable Y in V , such that there is a directed path from X to Y (i.e., a chain of directed edges leading from X to Y), then Y is called a 'descendant of X ' (Spirtes et al., 2000, p. 8–10).

A probabilistic causal model M is a formal structure consisting of a directed acyclic graph and a probability distribution over the variables in that graph (Gebharter, 2017, p. 357). One of the standard conditions imposed on graphs constituting causal models is the causal Markov condition:

Causal Markov condition: A causal model M satisfies the causal Markov condition iff for each variable X in the set of variables V of M : conditional on its parents in V , X is probabilistically independent of all other variables in V except its descendants (see e.g. Pearl, 2000, p. 30; Spirtes et al., 2000, p. 29; Woodward, 2003, p. 64).

The causal Markov condition implies that if there is a probabilistic dependence relation between X and Y conditional on the parents of X in V , then Y must be a descendant of X in the graph constituting M . However, the causal Markov condition can also be satisfied by models in which X is a descendant of Y , but there is no probabilistic dependence relation between X and Y . This is excluded if causal models are additionally required to satisfy the faithfulness condition:

Faithfulness condition: A causal model M satisfies the faithfulness condition iff all conditional independence relations between the variables in M are entailed by the causal Markov condition applied to M (Spirtes et al., 2000, p. 31).

The purpose of introducing these two conditions is to be able to give a causal interpretation to the directed edges that occur in a model. We will see below that the causal Markov condition and the faithfulness condition alone are not sufficient to guarantee that the probabilistic dependencies in a model can be mapped to causal dependencies. However, if we assume that the set of variables constituting a model M is chosen in a way that allows such a mapping, then the following relations hold: If X and Y are variables in M , then X is a direct cause of Y according to M iff there is a directed edge from X to Y with no intermediate variables in M 's graph.¹ Y is a direct or contributing cause of X according to M iff X is a descendant of Y in M 's graph.² Given the relationship between graphs and probabilities specified by the causal Markov condition, this in turn implies that if there is a probabilistic dependence relation between X and Y conditional on the parents of X in a model M , then either X is a direct or contributing cause of Y according to M , or Y is a direct or contributing cause of X according to M .

The notions of direct and contributing cause are characterized relative to a set of variables. The choice of the variables included in a causal model is key, since models may misrepresent causal structures if the variables are not chosen in an inadequate way. This is particularly relevant in cases where non-causal dependence relations between variables are brought about by common causes. The relation between dry cough and elevated temperature mentioned above is an example. Another paradigmatic example is the covariation between the occurrence of storms in a certain region

¹ Strictly speaking, the variables themselves do not stand in causal relations, but only represent entities standing in causal relations. However, as is common in the causal modeling literature, I will use expressions such as 'variable X is causally relevant to variable Y ' as shorthand for 'the entities represented by the values of X are causally relevant to the entities represented by the values of Y '.

² The last two conditions presuppose that causation is transitive, that is, that if X is causally relevant to Y and Y is causally relevant to Z , then X is also causally relevant to Z . This is a standard assumption in the causal modeling literature (see, e.g., Spirtes et al., 2000, p. 20). However, Hitchcock has argued that there may be failures of transitivity, and that these should be captured by causal models (Hitchcock, 2001, 2007). In what follows, I will leave that complication aside. The failures of transitivity to which Hitchcock refers are cases in which Y depends counterfactually/probabilistically on X , and Z depends counterfactually/probabilistically on Y , but Z does not depend counterfactually/probabilistically on X . In these cases, there is no probabilistic dependence relation between X and Z , although there is a directed path from X to Z . Such cases are not relevant to the present discussion, which focusses on how causal models can deal with cases where there are probabilistic dependence relations, but these do not indicate causal relations.

and the observation that the barometers in that region dropped quickly. Such scenarios can be described by the following three variables:

St: 1 if there is a storm in region R at time t_2 ; 0 otherwise.

B: difference of barometer readings between times t_1 and t_2 in region R.

A: difference of atmospheric pressure between times t_1 and t_2 in region R.

Barometer readings at different times are good predictors of storms: if barometers in region R drop strongly between t_1 and t_2 , this will increase the probability of a storm in region R at t_2 . However, the probabilistic dependence relation between St and B is not causal, but due to the presence of a common cause – differences of atmospheric pressure in region R, described by variable A. Accordingly, if the set of variables under consideration included only B and St, B might be misclassified as a cause of St or vice versa.

Therefore, causal models are usually required to satisfy the causal sufficiency condition in addition to satisfying the causal Markov and the faithfulness condition (Baumgartner, 2013, p. 9; Spirtes & Scheines, 2004, p. 836–837; Zhang & Spirtes, 2011, p. 337):

Causal sufficiency: A set of variables V is causally sufficient iff for any two variables X and Y contained in V and all variables Z : if Z is both a direct cause of X and a direct cause of Y , then Z is in V , too (Spirtes & Scheines, 2004, p. 836; Spirtes, 2010, p. 1651; Zhang & Spirtes, 2011, p. 337).³

The causal sufficiency condition is usually applied to sets of variables, rather than to causal models. However, I will also apply it to models and define a model as causally sufficient iff its set of variables is causally sufficient.

In the barometer-storm example, a model including only B and St is not causally sufficient, since it does not contain the common cause(s) of B and St. If we consider a causally sufficient model including the variable A in addition to B and St, the putative causal relation between B and St disappears: A screens off B from St, that is, $p(\text{St} | B \& A) = p(\text{St} | A)$.

³ There are two possible alternative ways to define the causal sufficiency condition. The first is due to Baumgartner, who defines a set of variables V as causally sufficient ‘iff any common cause C of two variables X and Y in V either belongs to V or has a cause that belongs to V or an effect that is located on all directed paths from C to X and from C to Y and that belongs to V ’ (Baumgartner, 2013, p. 9). The difference between this definition and the definition of causal sufficiency given above is that the common causes that must be included in V need not be direct common causes of variables included in V but could also be contributing causes. The second alternative definition of causal sufficiency is given by Spirtes, Glymour, and Scheines, who define a set of variables V as causally sufficient iff, whenever a variable X is a direct cause of two or more variables in V , then either X is in V or ‘the joint probability of all variables in V ... [is] the same on each value of X that occurs in the population’ (Spirtes et al., 2000, p. 22, fn. 1). In what follows, I will not consider these alternative formulations of causal sufficiency, since using them instead of the definition given above would not change my argument in any substantial way.

Prima facie, causal sufficiency is a powerful condition to impose on causal models. At the very least, it seems to guarantee that common cause structures are adequately captured and that there are no cases where non-causal relations are misrepresented as causal because the model contains too few variables. In the next section, however, I raise three objections to causal sufficiency.

3 Three problems for causal sufficiency

The first problem for causal sufficiency arises from the requirement that a causally sufficient model must include *all* variables that are direct common causes of two or more variables already included in the model. To see why this requirement is problematic, consider again the model that describes the relationship between barometer readings, atmospheric pressure differences, and storms. The atmospheric pressure differences could alternatively be described by the following two variables:

A1: atmospheric pressure in region R at time t1.

A2: atmospheric pressure in region R at time t2.

A1 and A2 are common causes of B and St. Therefore, if the model included A1 and A2 instead of A, B would also be screened off from St, that is, $p(\text{St} | B \& A1 \& A2) = p(\text{St} | A1 \& A2)$. The problem is, however, that if *all* common causes of B and St must be included in the model, then the model must include all three variables, A, A1, and A2. But then, to determine whether B is screened off from St by their common causes, one must determine whether $p(\text{St} | B \& A \& A1 \& A2) = p(\text{St} | A \& A1 \& A2)$. Since A is just the difference of A1 and A2, some combinations of the values of A, A1, and A2 are impossible, for instance, $A = 35.0$ hPa, $A1 = 1015.0$ hPa, and $A2 = 990.0$ hPa. This means that $p(A = 35.0 \text{ hPa} \& A1 = 1015.0 \text{ hPa} \& A2 = 990.0 \text{ hPa})$ assumes the value zero, and the corresponding conditional probabilities, for instance, $p(\text{St} = 1 | A = 35.0 \text{ hPa} \& A1 = 1015.0 \text{ hPa} \& A2 = 990.0 \text{ hPa})$, are not well defined.

This problem is usually avoided by requiring that the variables included in a causal model must not stand in non-causal necessary dependence relations to each other. Woodward formulates the so-called ‘independent fixability’ condition for his interventionist framework, according to which the set V of variables constituting an interventionist causal model must be such that every combination of the values of the variables in V must be metaphysically possible and not be excluded for logical, mathematical or conceptual reasons (Woodward, 2015, p. 316). This precludes variables such as A, A1, and A2 from occurring in the same causal model.

In probabilistic causal models, an even stronger condition is needed, because the problem that the joint probability of the values of certain variables might be equal to zero can also arise if the variables stand in deterministic dependence relations that hold only with nomological necessity, but not with metaphysical, logical, mathemati-

cal, or conceptual necessity.⁴ Variables appearing together in fundamental physical laws could be an example. It is therefore plausible to assume that the set of variables constituting a probabilistic causal model must satisfy the following condition (which is a slight adaptation of Woodward's independent fixability condition):

Independent fixability+: A set of variables V satisfies the independent fixability+ condition iff every combination of the values of the variables in V is nomologically (and thus metaphysically and logically) possible.⁵

The first problem with the causal sufficiency condition is that taken literally, it leads to violations of independent fixability+. It requires that *all* direct common causes of tuples of variables already included in the model must also be included. However, as the barometer-storm example illustrates, if X and Y are included in a model, then there may be several variables which are all direct common causes of X and Y and which stand in necessary dependence relations to each other. If all of these variables are included, the set of variables that makes up the resulting model violates independent fixability+.

One could try to refine the causal sufficiency condition to avoid this problem. If causal sufficiency does not require that all common causes of tuples of variables in a model M should be included in M , but only all common causes that can be included without violating the independent fixability+ condition, then the problem that the resulting model might contain undefined conditional probabilities disappears. However, I will leave open whether and how exactly the causal sufficiency condition can be reformulated, since it faces two other difficulties.

To see a second difficulty, reconsider the definition of causal sufficiency: a set of variables V is causally sufficient iff for any two variables X and Y contained in V and all variables Z : if Z is both a direct cause of X and a direct cause of Y , then Z is in V , too. How should we understand the notion that Z is a direct cause of X and of Y ? One option is to understand the notion that Z is a direct cause of X and of Y relative to a model consisting of a certain set of variables. Then, the crucial question is, what is the set of variables relative to which Z is a direct cause of X and a direct cause of Y ?

Suppose first that the relevant set of variables is V itself, that is, the set of variables that make up the model that is supposed to be causally sufficient. Then the causal sufficiency condition is trivially true. For a variable can be a direct cause of another variable relative to a set S only if it is a member of S . Therefore, the assumption that Z is a direct cause of X and Y relative to V directly implies that Z must be in V . It follows that the condition that Z is a direct cause of X and a direct cause of Y must be

⁴ Another reason why the variables constituting a model must not stand in deterministic relations to each other is that otherwise the model may violate the faithfulness condition (see Spirtes et al., 2000, p. 53–54).

⁵ There has been a debate in recent years about whether causal models can be *hybrid* or *mixed* in the sense that they contain not only standard causal relations, but also dependence relations of other types, especially supervenience or grounding relations. A number of authors have attempted to adapt the formalism of causal models to models which violate Woodward's independent fixability condition (e.g., Gebharter, 2017; Kroedel & Schulz, 2016; Stern & Eva, 2023; for discussion see also Kistler, 2013; Shapiro, 2010; Eronen & Brooks, 2014). In what follows, I will leave open whether and how my argument might apply to these other versions of the formalism of causal models.

understood as referring to a set of variables V^* that is distinct from V (see also Peters et al., 2017, p. 171–172, who assume V^* to be a superset of V).

Now consider such a set $V^* \neq V$. V^* must satisfy the adequacy conditions for causal models. Otherwise, the claim that Z is a direct cause of X and Y relative to V^* would not be justified. In particular, V^* must be causally sufficient (Peters et al., 2017, p. 172, also seem to assume this). But this leads to an infinite regress. V^* is causally sufficient iff for any two variables X and Y contained in V^* and all variables Z : if Z is both a direct cause of X and a direct cause of Y relative to some set V^{**} , then Z is in V^* . According to the argument given in the previous paragraph V^{**} must be distinct from V^* . However, V^{**} must also be causally sufficient. But then, the argument can be reiterated: whether V^{**} is causally sufficient can only be determined relative to some set V^{***} , which in turn must be distinct from V^{**} and causally sufficient, and so on. Therefore, if the notion that Z is a direct cause of X and of Y is understood relative to a model consisting of a certain set of variables, the causal sufficiency condition is either trivially satisfied, or leads to an infinite regress.

The remaining option is to understand the notion of direct common cause that appears in the causal sufficiency condition as an undefined primitive. Such a move, which would avoid both the triviality problem and the regress problem, could be justified by the observation that causal models must always presuppose certain notions, including causal notions, as primitive, and that this is not per se problematic. The causal modelling approach is usually not considered to be reductive, that is, its aim is not to reduce the notion of causation completely to non-causal notions. As long as the conclusion that X is causally relevant to Y is based only on assumptions about causal relations *other than* the one between X and Y , the approach is not problematically circular (for a related consideration, applied to the interventionist framework of causation, see Woodward, 2003, p. 104–16). Still, an approach that presupposes fewer primitives is *ceteris paribus* superior to one that presupposes more. I will come back to this point in the next section.⁶

The third problem with the causal sufficiency condition is that it is too weak to rule out a relevant type of problematic structure. As pointed out above, the reason for requiring causal models to be causally sufficient is that there may be probabilistic dependence relations that do not indicate causal relations but are due to common-cause structures. This type of structure can be adequately covered by the causal sufficiency condition. However, structures containing variables that are probabilistically related by accident (Williamson, 2004, p. 52) cannot be adequately covered.

Consider the following two variables, which describe Sober's famous example of the relationship between the sea level in Venice and bread prices in England:

SV: sea level in Venice.

BE: bread prices in England.

⁶ It should be noted that Peters, Janzing and Schölkopf also replace the causal sufficiency condition with a condition they call 'interventional sufficiency', which does not rely on the notion of direct common causes (Peters et al., 2017, p. 172).

Since both bread prices in England and the sea level in Venice have risen (more or less) steadily over the last two centuries, BE and SV are strongly positively correlated. However, there is no common causal process to account for this dependence relation. The two variables are correlated just because their values develop in parallel over time (see Sober, 2001, p. 331–332; further cases of variables that are probabilistically related without standing in direct or indirect causal relations to each other are discussed by Cartwright, 1989, p. 114–115; Spirtes et al., 2000, p. 32–38; Williamson, 2004, p. 52–57).

Now consider a model consisting of the set of variables $\{BE, SV\}$. This model is causally sufficient because BE and SV have no common causes. BE and SV also satisfy the independent fixability+condition, since they are not necessarily related. Moreover, if the causal Markov condition holds, as we assume, the probabilistic dependence between BE and SV entails that one must draw an edge between them. This is because if there is no edge between BE and SV, then BE and SV have no parents or descendants in the model. But then the causal Markov condition implies that they have to be probabilistically independent, which is not the case. It follows that the model constituted by $\{BE, SV\}$ misrepresents the true causal structure, because there must be an edge between BE and SV, even though the two variables are not causally related. Thus, the three conditions we have considered so far – the causal Markov condition, faithfulness, and causal sufficiency – are not sufficient to handle structures that contain correlations that are not due to causal relevance relations or common causes.⁷

An analogous problem occurs if we replace BE in the model with the following variable, so that the new model is constituted by the set of variables $\{PE, SV\}$:

PE: number of households living below the poverty line in England.

It is plausible to suppose that bread prices in England (BE) are causally relevant to PE, but the sea level in Venice (SV) is not. However, given the strong positive correlation between SV and BE, there is a strong positive correlation between SV and PE: the higher the sea level in Venice, the higher the number of households living below the poverty line in England. Therefore, a model that includes only SV and PE (but not BE) would also misrepresent the true causal structure, because it would have to include a directed edge between SV and PE, even though the two variables are causally unrelated.

⁷ One possible response is that causal models should satisfy the additional constraint that the error terms of the variables that constitute them should be probabilistically independent (Papineau, 2022). If the model constituted by $\{BE, SV\}$ does not satisfy this condition, it is inadequate from the start, in which case it is unproblematic if it misrepresents the true causal structure.

However, what the error terms of the variables in the model are depends on what the directed edges between the variables are (and on the structural equations describing the functional relations represented by these edges). If the model constituted by SV and BE contains a directed edge from BE to SV because of the strong covariation between SV and BE, it is likely to satisfy this condition. The error terms of SV relative to this model are variables describing influences on SV that are independent of the putative (that is, statistical) influence of BE. However, these are likely to be independent of the error terms of BE (since SV and BE do not have any common causes). At the very least, it cannot simply be taken for granted that the model including only SV and BE with a directed edge between them violates the condition of independent error terms.

In the next section, I argue that such structures are better covered by requiring causal models to be monotonic, and that if causal sufficiency is replaced by monotonicity, the causal modelling approach also fares better with respect to the other two problems described in this section.

4 Monotonicity

In a recent paper, Papineau mentions that the causal sufficiency condition may be circular and hints at a solution to this problem:

A revised reductive suggestion would now be that causal relations are nothing over and above those patterns of correlation that imply them ... in any causally sufficient set of variables. (Would not the need to specify causal sufficiency here render this suggestion inadmissibly circular as a reduction of causation? But this specification can be finessed away. We can simply say causal relations are nothing over and above the patterns of correlation that imply them in sets of variables whose verdicts are not overturned by the inclusion of further variables.) (Papineau, 2022, p. 253).

Papineau's point seems to be that a reductive analysis of causation should not rely on the causal sufficiency condition, because causal sufficiency presupposes the notion of causation, and this makes the analysis circular. As pointed out in the previous section, probabilistic causal models are usually not intended to be reductive. Therefore, the observation that causal sufficiency presupposes causal notions does not per se render this approach problematic. However, it was also pointed out in the previous section that there are two other problems with causal sufficiency. According to Papineau, the requirement that the (putative) causal relations in a model must not disappear when common causes are added should be replaced by the requirement that the (putative) causal relations in a model must not disappear when additional variables – causally related to the variables already included in the model or not – are added. In the context of probabilistic causal models, this idea can be used as the basis for the following monotonicity condition, which, as I will argue, is superior to causal sufficiency with respect to all three problems identified in the previous section:

Monotonicity: A model M consisting of a set of variables V is monotonic iff for any X and Y in V : if X is a direct or a contributing cause of Y according to M , then X would still be a direct or a contributing cause of Y according to any M' consisting of a set of variables V' , such that (i) $V \subset V'$, and (ii) the variables in V' satisfy the independent fixability+ condition.⁸

⁸ The monotonicity condition is similar to the expandability condition that Papineau introduces into his formal system:

Let us say that a system S of equations with exogenous independence is expandable if, for any further variables correlated with those in S , there is a larger system of equations covering those further variables that also satisfies exogenous independence and which has S as a subsystem. (Papineau, 2022, p. 264)

There are two things to note about this definition. The first is that it still presupposes that both the causal Markov condition and the faithfulness condition hold. The second thing to note is the function of the phrase ‘direct or contributing cause’, which occurs twice. The notion of direct causation is clearly not monotonic (Woodward, 2008, p. 209; see also Parkkinen, 2022, p. 192–194). Often it is a matter of convention or pragmatic considerations how many variables on a causal path are included in a set of variables. For instance, there could be a model including a variable that describes whether somebody throws a paper ball and another variable that describes whether the ball lands in a basket. If no other variables are included in the model, the first variable is a direct cause of the second. If one adds further variables lying on the same causal path, for instance, variables describing the position and momentum of the ball on its way to the basket, the throw is not a direct cause of the paper ball’s landing in the basket, but still a contributing cause. Such cases are covered by the monotonicity condition. The condition requires that if X is a direct cause of Y with respect to a set of variables V , then X remains a direct *or a contributing cause* of Y relative to any extended set of variables satisfying condition (ii) (for a very similar consideration, see Woodward, 2008, p. 209). If X is a contributing cause of Y , then X remains a contributing cause of Y under suitable extensions of the set of variables. However, requiring that X remains a direct cause of Y under suitable extensions of the set of variables would be too strong.

Condition (ii), that the variables hypothetically added to the model must be such that the extended set of variables does not violate the independent fixability+ condition, guarantees that the monotonicity condition does not face the first problem for causal sufficiency identified in the previous section. If the extended set of variables V' satisfies the independent fixability+ condition, then there will be no cases where the conditional probabilities over the variables in V' are undefined because some variables in V' stand in necessary dependence relations to each other. Condition (ii) thus ensures that the hypothetical extensions of the models under consideration can be covered by the standard formalism of probabilistic causal models.

Moreover, the monotonicity condition avoids the other two problems. To see how it avoids the third problem, consider again a model M consisting of the set of variables $V = \{SV, PE\}$ (sea level in Venice, households living in poverty in England). As pointed out in the previous section, SV and PE stand in a probabilistic dependence relation to each other: a higher sea level in Venice increases the probability that there will be more households living below the poverty line in England. Therefore, the causal Markov condition implies that there must be a directed edge between SV and PE , and the model misrepresents the true causal structure.

However, there is a crucial difference between the monotonicity condition and Papineau’s notion of expandability. The latter requires that the complete causal structure of the system remains intact under the addition of new variables (assuming independence of the exogenous variables is still given). This condition implies that if variables X and Y are already included in S and if there is *no* causal relation between X and Y according to S , then expanding S should not create a causal relation between X and Y . However, several authors have argued that this condition is too strong (Spirtes et al., 2000, p. 21; Statham, 2018, Statham’s argument refers to the context of interventionist causal models, but could easily be adapted to apply to probabilistic causal models as well). The monotonicity condition avoids this difficulty.

However, M violates the monotonicity condition. Consider a model M' which is constituted by the set of variables $V' = \{SV, PE, BE\}$ (where BE describes the bread prices in England), and which contains a directed edge from BE to PE . V' is a superset of V (condition (i)), and the values of the variables in V' satisfy the independent fixability+condition (condition (ii)). But BE , the variable added to V , is a parent of PE and screens off SV from PE : $p(PE \mid SV \ \& \ BE) = p(PE \mid BE)$. Accordingly, the faithfulness condition implies that there is no directed edge between SV and PE , and the putative causal relation between SV and PE disappears. It follows that there is a model M' whose set of variables satisfies conditions (i) and (ii), but which does not contain a directed edge from SV to PE . Therefore, the initial model M violates the monotonicity condition.

Analogous reasoning applies to Sober's original case, that is, models containing only the variables SV and BE . SV and BE have no *common* causes, but they are certainly not uncaused. Suppose that X_1, \dots, X_n are the causes of SV , such as vertical land movement and melting of the Arctic ice shield (Zanchettin et al., 2021), and that SV depends probabilistically on each of the X_i s. Suppose further that X_1, \dots, X_n are unconditionally independent.

Now consider a model M consisting of the set of variables $V = \{SV, BE\}$ and containing a directed edge from SV to BE . To see that M is not monotonic, consider an extended model M' consisting of the set of variables $V' = \{SV, X_1, \dots, X_n, BE\}$ and containing directed edges from each of the X_i s to SV , but not between the X_i s (because they are unconditionally independent). V' is a superset of V (condition (i)) and satisfies the independent fixability+condition (condition (ii)). Furthermore, X_1, \dots, X_n screen off BE from SV : $p(SV \mid BE \ \& \ X_1 \ \& \ \dots \ \& \ X_n) = p(SV \mid X_1 \ \& \ \dots \ \& \ X_n)$. Again, this shows that the original model M violates the monotonicity condition: the set of variables constituting the extended model M' satisfies conditions (i) and (ii), but does not contain a directed edge from SV to BE (because of the faithfulness condition). Analogous reasoning applies if the edge between SV and BE goes in the opposite direction, that is, from BE to SV : in this case, the probabilistic dependence relation that holds between SV and BE can be screened off by variables representing the causes of BE .

An immediate objection here is that the extended models are not adequate either, because they also violate the monotonicity condition. The model consisting of the set of variables $V = \{SV, PE\}$ is not monotonic, because in the model consisting of the extended set of variables $V' = \{SV, PE, BE\}$, the probabilistic dependence between SV and PE is screened off by BE . However, the model consisting of $V' = \{SV, PE, BE\}$ does *not* satisfy the monotonicity condition either, since (as we saw above) the probabilistic dependence relation between SV and BE would be screened off if the variables X_1, \dots, X_n , describing the causes of SV , were added to it. And this extended model is not monotonic either, because BE depends probabilistically on X_1, \dots, X_n , and this probabilistic dependence relation would be screened off if variables describing the causes of BE (such as higher flour prices and rising production costs) were added. And arguably, even this extended model is not monotonic because there will be statistical dependence relations between (some of) the causes of SV (i.e., X_1, \dots, X_n) and (some of) the causes of BE , which would disappear if further causes were added to the model.

In this case, however, what looks like a regress structure is not problematic. The hypothetical model constituted by V' used to show that the model constituted by $V = \{SV, PE\}$ is not monotonic must satisfy conditions (i) and (ii), but these conditions do not involve monotonicity and do not require primitively given knowledge of the true causal relations between the variables in V' . It is possible that M' misrepresents the true causal structure, and in a further step, one might consider M' and find that it is not monotonic either. But this does not mean that M' cannot be used to determine whether the original model M is monotonic.

This consideration also shows that monotonicity is superior to causal sufficiency with respect to the second problem described in the previous section. According to the monotonicity condition, the variables hypothetically added to V need not satisfy any constraints other than the independent fixability+condition, and the extended model M' need not be monotonic (let alone causally sufficient). Therefore, there is no threat of regress. Moreover, as the monotonicity condition does not presuppose that the variables hypothetically added to V are (direct or indirect) causes of variables already contained in V , it assumes fewer causal relations as primitively given than the causal sufficiency condition.

The monotonicity condition captures scientific practice in the following sense: if observational data show that there is a correlation between variables, then one should try to determine whether one of the variables is indeed causally dependent on the other, or whether the dependence is due to other factors that can be described by additional variables. These additional variables may represent common causes of the correlated variables, but one should always be open to considering other potentially relevant factors. Obviously, there are an infinite number of factors to consider, and epistemically, one can never be completely sure that a model is monotonic. Monotonicity should therefore be understood as an ideal that scientists developing causal models should strive to approximate.

5 A note on interventionism

A possible objection to the argument of this paper is that the problem that is solved by imposing monotonicity as an additional requirement on causal models has already been solved by Woodward's interventionist theory of causation. According to this version of the causal modeling approach, a variable X included in a causal model is causally relevant to a variable Y occurring in the same model iff there is an intervention on the value of X that changes the value or the probability distribution of Y , provided that the values of all other variables in the model that are not on the causal path between X and Y are held fixed by interventions (Woodward, 2003; Hitchcock, 2001, 2007).

The crucial difference between Woodward's approach and the framework discussed so far is that Woodward's approach requires the notion of an intervention as an additional component. Interventions are characterized by intervention variables, which are defined as follows: 'I is an intervention variable for X , with respect to Y , if it meets the following conditions:

- (1) I is causally relevant to X.
- (2) I is not causally relevant to Y through a route that excludes X.
- (3) I is not correlated with any variable Z that is causally relevant to Y through a route that excludes X, be the correlation due to I's being causally relevant to Z, Z's being causally relevant to I, I and Z sharing a common cause, or some other reason.
- (4) I acts as a switch for other variables that are causally relevant to X. That is, certain values of I are such that when I attains those values, X ceases to depend upon the values of other variables that are causally relevant to X.' (Woodward & Hitchcock, 2003, p. 12–13).

According to Woodward's approach, the variable BE (representing bread prices in England) is causally relevant to the variable SV (representing the sea level in Venice) iff there is a possible intervention I on BE with respect to SV that changes the probability distribution of SV. Since I must satisfy the conditions of an intervention variable for BE, it must in particular be independent of all variables that are causally relevant to SV through a route that excludes BE (condition (3)). It is plausible to assume that if I satisfies this condition, the change in the value of BE will have no impact on the probability distribution of SV. For example, the bread prices in England could possibly be lowered by paying government subsidies to bakeries. However, such an intervention would not affect the sea level in Venice.

In general, since the causal chains leading to BE and SV do not overlap, interventions on one of these variables that satisfy condition (3) will not affect the causes of the other variable and thus will not change the probability distribution of the other variable. Accordingly, Woodward's interventionist criterion of causation correctly implies that there is no causal relation between BE and SC.

Condition (3) of Woodward's definition of an intervention can thus be seen as the interventionist solution to the problem of coincidental probabilistic dependence relations. However, there are two significant differences between the monotonicity condition and Woodward's approach. The first is that Woodward's solution only works if one is willing to accept his interventionist framework, and in particular the notion of hypothetical interventions. In a framework that relies only on observational probabilistic relations, such as the one described above, this solution is not available.⁹

The second difference is that Woodward's solution requires more causal information than the monotonicity condition. The definition of an intervention is not relativized to a set of variables. Thus, condition (3) requires that in order to determine whether I is an intervention variable for X with respect to Y, one must know *all* the causes of Z (not just those included in the model under consideration). The monotonicity condition, on the other hand, does not require any knowledge of the causal relations between the variables hypothetically added to the model and the variables already included in the model.

⁹ Baumgartner and Falk develop a formal regularity theory of causation (based on INUS conditions) and require causal structures to satisfy a permanence condition that is structurally similar to the monotonicity condition that I propose (Baumgartner & Falk, 2023). A detailed discussion of the connections between their approach and the one defended in this paper would go beyond the scope of the present argument.

Ultimately, the solution chosen to solve the problem of accidental probabilistic relations depends on the underlying framework and the additional commitments one is willing to make. However, the argument of this paper shows that, at least in the context of probabilistic causal models, replacing the causal sufficiency condition with the monotonicity condition has several advantages. Monotonicity avoids the conflict with the requirement that the variables included in a causal model must not stand in non-causal deterministic relations to each other. It requires fewer assumptions about which causal relations are primitively given than causal sufficiency. Finally – and this is perhaps the most systematically relevant consequence – it allows the causal modelling approach to deal better with the problem of accidental probabilistic relations.

Acknowledgements I would like to thank the editors, Stephan Krämer, Stefan Roski, and Singa Behrens, for inviting me to contribute to this special issue on *Difference-Making and Explanatory Relevance* and for organizing a very constructive and high-quality review process for my paper. Two anonymous referees provided incredibly detailed, patient, and competent comments on my paper over several rounds of revisions, and clearly helped me to bring my argument to its present – in my opinion much improved – form. Last but not least, I would like to thank Thomas Blanchard, Nicole Dolby-Rathgeb, Alexander Gebharder, Matthias Rolffs, and Jonas Werner for extremely valuable input on various parts of my argument.

Funding information The research leading to the results presented in this paper received funding from the Swiss National Science Foundation (SNSF) under Grant Agreement No 100019E_189589/1. Open access funding provided by University of Bern

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baumgartner, M. (2013). Rendering interventionism and non-reductive physicalism compatible. *Dialectica*, 67(1), 1–27. <https://doi.org/10.1111/1746-8361.12008>.
- Baumgartner, M., & Falk, C. (2023). Boolean difference-making: A modern regularity theory of causation. *British Journal for the Philosophy of Science*, 74(1), 171–197. <https://doi.org/10.1093/bjps/axz047>.
- Cartwright, N. (1989). *Nature's capacities and their measurement*. Oxford University Press.
- Eronen, M. I., & Brooks, D. S. (2014). Interventionism and supervenience: A new problem and provisional solution. *International Studies in the Philosophy of Science*, 28(2), 185–202. <https://doi.org/10.1080/02698595.2014.932529>.
- Gebharder, A. (2017). Causal exclusion and causal bayes nets. *Philosophy and Phenomenological Research*, 95(2), 353–375. <https://doi.org/10.1111/phpr.12247>.
- Hitchcock, C. (2001). The intransitivity of causation revealed in equations and graphs. *Journal of Philosophy*, 98(6), 273–299. <https://doi.org/10.2307/2678432>.
- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *Philosophical Review*, 116(4), 495–532. <https://doi.org/10.1215/00318108-2007-012>.
- Hitchcock, C. (2023). Causal models. In E.N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/spr2023/entries/causal-models/>.

- Kistler, M. (2013). The interventionist account of causation and non-causal association laws. *Erkenntnis*, 78(1), 1–20. <https://doi.org/10.1007/s10670-013-9437-4>.
- Kroedel, T., & Schulz, M. (2016). Grounding mental causation. *Synthese*, 193(6), 1909–1923. <https://doi.org/10.1007/s11229-015-0820-3>.
- Papineau, D. (2022). The statistical nature of causation. *The Monist*, 105(2), 247–275. <https://doi.org/10.1093/monist/onab034>.
- Parkkinen, V. P. (2022). Variable relativity of causation is good. *Synthese*, 200(3), 1–21. <https://doi.org/10.1007/s11229-022-03676-0>.
- Pearl, J. (2000). *Causality: Models, reasoning, and inference*. Cambridge University Press.
- Peters, J., Janzing, D., & Schölkopf, B. (2017). *Elements of causal inference. Foundations and learning algorithms*. The MIT Press.
- Shapiro, L. (2010). Lessons from causal exclusion. *Philosophy and Phenomenological Research*, 81(3), 594–604. <https://doi.org/10.1111/j.1933-1592.2010.00382.x>.
- Sober, E. (2001). Venetian sea levels, British bread prices, and the principle of the common cause. *British Journal for the Philosophy of Science*, 52(2), 331–346. <https://doi.org/10.1093/bjps/52.2.331>.
- Spirtes, P. (2010). Introduction to causal inference. *Journal of Machine Learning Research*, 11, 1643–1662.
- Spirtes, P., & Scheines, R. (2004). Causal inference of ambiguous manipulations. *Philosophy of Science*, 71(5), 833–845. <https://doi.org/10.1086/425058>.
- Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, prediction and search* (2nd ed.). The MIT Press.
- Statham, G. (2018). Woodward and variable relativity. *Philosophical Studies*, 175(4), 885–902. <https://doi.org/10.1007/s11098-017-0897-2>.
- Stern, R., & Eva, B. (2023). Antireductionist interventionism. *British Journal for the Philosophy of Science*, 74(1), 241–267. <https://doi.org/10.1086/714792>.
- Williamson, J. (2004). *Bayesian nets and causality: Philosophical and computational foundations*. Oxford University Press.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford University Press.
- Woodward, J. (2008). Response to Strevens. *Philosophy and Phenomenological Research*, 77(1), 193–212. <https://doi.org/10.1111/j.1933-1592.2008.00181.x>.
- Woodward, J. (2015). Interventionism and causal exclusion. *Philosophy and Phenomenological Research*, 91(2), 303–347. <https://doi.org/10.1111/phpr.12095>.
- Woodward, J., & Hitchcock, C. (2003). Explanatory generalizations, part I: A counterfactual account. *Noûs*, 37(1), 1–24. <https://doi.org/10.1111/1468-0068.00426>.
- Zanchettin, D., Bruni, S., Raicich, F., Lionello, P., Adloff, F., Androssov, A., Antonioli, F., Artale, V., Carminati, E., Ferrarin, C., Fofonova, V., Nicholls, R. J., Rubinetti, S., Rubino, A., Sannino, G., Spada, G., Thiéblemont, R., Tsimplis, M., Umgiesser, G., Vignudelli, S., Wöppelmann, G., & Zerbini, S. (2021). Sea-level rise in Venice: Historic and future trends (review article). *Natural Hazards and Earth System Sciences*, 21(8), 2643–2678. <https://doi.org/10.5194/nhess-21-2643-2021>.
- Zhang, J., & Spirtes, P. (2011). Intervention, determinism, and the causal minimality condition. *Synthese*, 182(3), 335–347. <https://doi.org/10.1007/s11229-010-9751-1>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.