



UvA-DARE (Digital Academic Repository)

Thinking About Causation: A Causal Language with Epistemic Operators

Barbero, F.; Schulz, K.; Smets, S.; Velázquez-Quesada, F.R.; Xie, K.

DOI

[10.1007/978-3-030-65840-3_2](https://doi.org/10.1007/978-3-030-65840-3_2)
[10.48550/arXiv.2010.16217](https://doi.org/10.48550/arXiv.2010.16217)

Publication date

2020

Document Version

Final published version

Published in

Dynamic Logic : New Trends and Applications

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Barbero, F., Schulz, K., Smets, S., Velázquez-Quesada, F. R., & Xie, K. (2020). Thinking About Causation: A Causal Language with Epistemic Operators. In M. A. Martins, & I. Sedlár (Eds.), *Dynamic Logic : New Trends and Applications: Third International Workshop, DaLí 2020, Prague, Czech Republic, October 9–10, 2020 : revised selected papers* (pp. 17-32). (Lecture Notes in Computer Science; Vol. 12569). Springer. https://doi.org/10.1007/978-3-030-65840-3_2, <https://doi.org/10.48550/arXiv.2010.16217>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)



Thinking About Causation: A Causal Language with Epistemic Operators

Fausto Barbero¹, Katrin Schulz², Sonja Smets^{2,3},
Fernando R. Velázquez-Quesada^{2(✉)}, and Kaibo Xie²

¹ University of Helsinki, Helsinki, Finland
fausto.barbero@helsinki.fi

² ILLC, Universiteit van Amsterdam, Amsterdam, The Netherlands
{K.Schulz, S.J.L.Smets, F.R.VelazquezQuesada, K.Xie}@uva.nl

³ Department of Information Science and Media Studies, University of Bergen, Bergen, Norway

Abstract. In this paper we propose a formal framework for modeling the interaction of causal and (qualitative) epistemic reasoning. To this purpose, we extend the notion of a causal model [11, 16, 17, 26] with a representation of the epistemic state of an agent. On the side of the object language, we add operators to express knowledge and the act of observing new information. We provide a sound and complete axiomatization of the logic, and discuss the relation of this framework to causal team semantics.

Keywords: Causal reasoning · Epistemic reasoning · Counterfactuals · Team semantics · Dependence

1 Introduction

In recent years a lot of effort has been put in the development of formal models of causal reasoning. A central motivation behind this is the importance of causal reasoning for AI. Making computers take into account causal information is currently one of the central challenges of AI research [9, 27]. There has also been tremendous progress in this direction after the earlier groundbreaking work in [23] and [28]. Advanced formal and computational tools have been developed for modelling causal reasoning and learning causal information, with applications in many different scientific areas. In this paper we want to extend this work further. The direction we want to explore is that of developing formal models of the interaction between causal and epistemic reasoning.

Even though the standard logical approach to causal reasoning [17, 18, 23] can model epistemic uncertainty¹, it does not permit reasoning about the interaction between causal and epistemic reasoning in the object language. Although recently there have been proposals adding probabilistic expressions to the object language (e.g., [21]), very little has been done on combining causal and qualitative epistemic reasoning.²

¹ E.g., by adding a probability distribution over a causal model's exogenous variables.

² See [5] for an exception, though the epistemic element is not made fully explicit in the language. We will come back to this approach in Sect. 6.

However, this kind of reasoning occurs frequently in our daily life, especially in connection with counterfactual thinking.

Example 1. *In front of Billie there is a button, which is connected to a circuit breaker and a sprinkler. If the circuit is closed, the sprinkler works if and only if the button is pushed. If the circuit is not closed, the sprinkler won't work, independently of the state of the button. Billie knows these causal laws. She can also see the button and the sprinkler, but she does not know the state of the circuit breaker. Suppose that at the moment the circuit is closed and the button is not pushed; as a result, the sprinkler is not working.*

We want to derive that Billie is not sure that if the button had been pushed, the sprinkler would have been working. Thus, we want to make inferences involving epistemic attitudes towards counterfactuals, which in turn explore causal dependencies. We also want to reason counterfactually about such epistemic attitudes. For instance, in the example, we also want to infer that if Billie had pushed the button and saw that the sprinkler works, then she would have known that the circuit is closed. To formalize this type of reasoning, we need a framework that combines causal reasoning with a model of epistemic attitudes.

Given the vast literature on epistemic logic, there is a lot of work that we can build on. This paper makes a start on combining the standard approach to causal reasoning [17, 18, 23] with tools from Dynamic Epistemic Logic (DEL; [4, 8, 13]). The main motivation for this choice is the dynamic character of both systems, even though this aspect will not be explored at depth here. For now we will only consider a very simple extension of the standard system of causal reasoning. But, as we will show, this basic extension already allows us to formalise some interesting concepts and formulate concrete questions for further research.

Outline. Section 2 introduces the standard approach to causal reasoning, and then Sect. 3 motivates in more detail the extension proposed here. Section 4 extends the standard causal modeling with means to express knowledge and external communication, and Sect. 5 provides a sound and complete axiomatization for the new system. Section 6 concludes the paper discussing the relationship with Causal Team Semantics [5, 6]. For space reasons, some proofs have been omitted; they can be found in the paper's full version.³

2 The Standard Causal Modelling Approach

What we refer to as the standard logic of causal reasoning was presented on [24], extended in [16], and then further developed in, among others, [11, 17, 25]. This section recall briefly the most important concepts and tools. The starting point is a formal representation of causal dependencies. This is done through *causal models*, which represent the causal relationships between a finite set of variables. These variables as well as their ranges of values are given by a *signature*. Throughout this text, let $\mathcal{S} = \langle \mathcal{U}, \mathcal{V}, \mathcal{R} \rangle$ be the *finite signature* where

³ arXiv:2010.16217 [cs.AI].

- $\mathcal{U} = \{U_1, \dots, U_m\}$ is a finite set of *exogenous* variables (those whose value is causally independent from the value of every other variable in the system),
- $\mathcal{V} = \{V_1, \dots, V_n\}$ is a finite set of *endogenous* variables (those whose value is completely determined by the value of other variables in the system), and
- $\mathcal{R}(X)$ is the finite non-empty range of the variable $X \in \mathcal{U} \cup \mathcal{V}$.⁴

A causal model is formally defined as follows.

Definition 1 (Causal model). A causal model is a triple $\langle \mathcal{S}, \mathcal{F}, \mathcal{A} \rangle$ where

- $\mathcal{S} = \langle \mathcal{U}, \mathcal{V}, \mathcal{R} \rangle$ is the model's signature,
- $\mathcal{F} = \{f_{V_j} \mid V_j \in \mathcal{V}\}$ assigns, to each endogenous variable V_j , a map

$$f_{V_j} : \mathcal{R}(U_1, \dots, U_m, V_1, \dots, V_{j-1}, V_{j+1}, \dots, V_n) \rightarrow \mathcal{R}(V_j).$$

The map f_{V_j} is sometimes called V_j 's structural function, and the set \mathcal{F} is called a set of structural functions for \mathcal{V} .

- \mathcal{A} is the valuation function, assigning to every $X \in \mathcal{U} \cup \mathcal{V}$ a value $\mathcal{A}(X) \in \mathcal{R}(X)$. For each endogenous variable, the valuation should comply with the variable's structural function. In other words, for every $V_j \in \mathcal{V}$, the following should hold:

$$\mathcal{A}(V_j) = f_{V_j}(\mathcal{A}(U_1), \dots, \mathcal{A}(U_m), \mathcal{A}(V_1), \dots, \mathcal{A}(V_{j-1}), \mathcal{A}(V_{j+1}), \dots, \mathcal{A}(V_n)).$$

In a causal model $\langle \mathcal{S}, \mathcal{F}, \mathcal{A} \rangle$, the functions in \mathcal{F} describe the causal relationship between the variables. Using these functional dependencies, we can define what it means for a variable to directly causally affect another variable.⁵

Definition 2 (Causal dependency). Let \mathcal{F} be a set of structural functions for \mathcal{V} . Given an endogenous variable $V_j \in \mathcal{V}$, rename each other variable in \mathcal{S} , the variables $U_1, \dots, U_m, V_1, \dots, V_{j-1}, V_{j+1}, \dots, V_n$, as X_1, \dots, X_{m+n-1} , respectively.

We say that, under the structural functions in \mathcal{F} , an endogenous variable $V_j \in \mathcal{V}$ is directly causally affected by a variable $X_i \in (\mathcal{U} \cup \mathcal{V}) \setminus \{V_j\}$ (in symbols, $X_i \hookrightarrow_{\mathcal{F}} V_j$) if and only if there is a tuple

$$(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{m+n-1}) \in \mathcal{R}(X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_{m+n-1})$$

and there are $x'_i \neq x''_i \in \mathcal{R}(X_i)$ such that

$$f_{V_j}(x_1, \dots, x'_i, \dots, x_{m+n-1}) \neq f_{V_j}(x_1, \dots, x''_i, \dots, x_{m+n-1}).$$

When $X_i \hookrightarrow_{\mathcal{F}} V_j$, we will also say that X_i is a causal parent of V_j . The relation $\hookrightarrow_{\mathcal{F}}^+$ is the transitive closure of $\hookrightarrow_{\mathcal{F}}$.

⁴ Given $(X_1, \dots, X_k) \in (\mathcal{U} \cup \mathcal{V})^k$, abbreviate $\mathcal{R}(X_1) \times \dots \times \mathcal{R}(X_k)$ as $\mathcal{R}(X_1, \dots, X_k)$.

⁵ This notion of a *direct cause* is adopted from [16]; it is related to the notion of a variable having a *direct effect* on another, as discussed in [23] in the context of Causal Bayes Nets. The notions defined here differ from Halpern's notion of *affect* [17], and this affects the axiomatization: axiom HP6 (Table 1) has the same function as C6 in [17] (ensuring that the canonical model is recursive), but does so in a slightly different way.

As it is common in the literature, we restrict ourselves to causal models in which circular causal dependencies do not occur.⁶

Definition 3 (Recursive causal model). A set of structural functions \mathcal{F} is recursive if and only if $\hookrightarrow_{\mathcal{F}}^+$ is a strict partial order (i.e., an asymmetric [hence irreflexive] and transitive relation, so there are no cycles). A causal model $\langle \mathcal{S}, \mathcal{F}, \mathcal{A} \rangle$ is recursive if and only if \mathcal{F} is recursive. In this text, a recursive causal model will be called simply a causal model.

The most important notion of this formalisation of causal reasoning is that of an *intervention*. This notion refers to the action of changing the values of variables in the system. Before we define an intervention formally, let us first introduce the notion of assignment.

Definition 4 (Assignment). Let $\mathcal{S} = \langle \mathcal{U}, \mathcal{V}, \mathcal{R} \rangle$ be a signature. An assignment on \mathcal{S} is an expression $\vec{X} = \vec{x}$ where \vec{X} is a tuple of different variables in $\mathcal{U} \cup \mathcal{V}$ (that is, $\vec{X} = (X_1, \dots, X_k) \in (\mathcal{U} \cup \mathcal{V})^k$ for some $k \in \mathbb{N}$, with $X_i \neq X_j$ for $i \neq j$), and $\vec{x} \in \mathcal{R}(\vec{X})$.

Now, an intervention that sets a variable X to the value x can be defined as an operation that maps a given model M to a new model $M_{X=x}$, which is the same except that the function determining the value of X is replaced by the constant function mapping X to x . In other words, X is cut off from all its causal dependencies and fixed to the value x .

Definition 5 (Intervention). Let $M = \langle \mathcal{S}, \mathcal{F}, \mathcal{A} \rangle$ be a causal model; let $\vec{X} = \vec{x}$ be an assignment on \mathcal{S} . The causal model $M_{\vec{X}=\vec{x}}^{\mathcal{F}} = \langle \mathcal{S}, \mathcal{F}_{\vec{X}=\vec{x}}^{\mathcal{F}}, \mathcal{A}_{\vec{X}=\vec{x}}^{\mathcal{F}} \rangle$, resulting from an intervention setting the values of variables in \vec{X} to \vec{x} , is such that

- $\mathcal{F}_{\vec{X}=\vec{x}}^{\mathcal{F}}$ is as \mathcal{F} except that, for each endogenous variable X_i in \vec{X} , the function f_{X_i} is replaced by a constant function f'_{X_i} that returns the value x_i regardless of the values of all other variables.
- $\mathcal{A}_{\vec{X}=\vec{x}}^{\mathcal{F}}$ is the unique valuation where (i) the value of each exogenous variable not in \vec{X} is exactly as in \mathcal{A} , (ii) the value of each exogenous variable X_i in \vec{X} is the provided x_i , and (iii) the value of each endogenous variable complies with its new structural function (that in $\mathcal{F}_{\vec{X}=\vec{x}}^{\mathcal{F}}$).⁷

Building on the notion of intervention, we can now extend a propositional language with a new type of sentence. The expression $[\vec{X}=\vec{x}]\gamma$ should be read as the counterfactual conditional *if the variables in \vec{X} were set to the values \vec{x} , respectively, then γ would be the case*.

⁶ The reason behind this restriction is that only acyclic relations are thought to have a causal interpretation (see [29] for an argument). The counterfactuals satisfy different logical laws if cyclic dependencies are allowed (see [17]).

⁷ Note that, since \mathcal{F} is recursive, the valuation $\mathcal{A}_{\vec{X}=\vec{x}}^{\mathcal{F}}$ is uniquely determined. First, the value of every exogenous variable U is uniquely determined, either from \vec{x} (if U occurs in \vec{X}) or else from \mathcal{A} (if U does not occur in \vec{X}). Second, the value of every endogenous variable V is also uniquely determined, either from \vec{x} (if V occurs in \vec{X} , as V 's new structural function is a constant) or else from the (recall: recursive) structural functions in $\mathcal{F}_{\vec{X}=\vec{x}}^{\mathcal{F}}$ (if V does not occur in \vec{X}).

Definition 6. Formulas ϕ of the language \mathcal{L}_C based on the signature \mathcal{S} are given by

$$\begin{aligned} \gamma &::= Z = z \mid \neg\gamma \mid \gamma \wedge \gamma && \text{for } Z \in \mathcal{U} \cup \mathcal{V} \text{ and } z \in \mathcal{R}(Z) \\ \phi &::= Z = z \mid \neg\phi \mid \phi \wedge \phi \mid [\vec{X}=\vec{x}]\gamma && \text{for } \vec{X}=\vec{x} \text{ an assignment on } \mathcal{S} \end{aligned}$$

The language makes free use of Boolean operators, but it forbids the nesting of intervention operators $[\vec{X}=\vec{x}]$ (see [11] for a way to extend the system with nested interventions). Formulas of \mathcal{L}_C are evaluated in causal models $\langle \mathcal{S}, \mathcal{F}, \mathcal{A} \rangle$. The semantic interpretation for Boolean operators is the usual; for the rest,

$$\begin{aligned} \langle \mathcal{S}, \mathcal{F}, \mathcal{A} \rangle \models Z=z & \quad \text{iff} \quad \mathcal{A}(Z) = z \\ \langle \mathcal{S}, \mathcal{F}, \mathcal{A} \rangle \models [\vec{X}=\vec{x}]\gamma & \quad \text{iff} \quad \langle \mathcal{S}, \mathcal{F}_{\vec{X}=\vec{x}}, \mathcal{A}_{\vec{X}=\vec{x}}^{\mathcal{F}} \rangle \models \gamma \end{aligned}$$

3 Limitations of the Standard System

The notion of a causal model contains an incredible amount of extra information compared to classical models. Not only does it tell us which variables depend causally on which other variables, but it also determines the exact character of this dependence. On the side of the language this wealth of information is then explored in terms of counterfactual conditionals using the concept of an intervention. This is where the actual causal reasoning happens. The standard logic of causal reasoning is in fact a logic of counterfactual reasoning. This is no accident: Judea Pearl, founder of the approach to causal reasoning introduced above, sees both concepts as intimately related. He argues that only when an agent can evaluate counterfactual conditionals does she fully engage with causal reasoning [26,27]. Counterfactual reasoning *is* the highest level of causal reasoning – a level that even the most advanced AI technology doesn't reach.⁸

Still, the basic causal framework has some limitations. An important one is that causal (or counterfactual) reasoning does not stand on its own: it does interact with other forms of reasoning. For instance, and as we illustrated in the introduction, counterfactual reasoning also considers the effect interventions have on the epistemic state of (observing) agents. We can reason that *If Peter had pushed the button, he would have known that his flashlight is broken*, which involves thinking about Peter's epistemic state after observing a causal intervention. This type of reasoning allows us to plan our actions (try out a flashlight before we take it for a night walk), and also influences our interaction with other agents (if you want Peter to come back from his walk, you should tell him to test his flashlight before he leaves). Therefore, a full account of the logic of causal reasoning needs to model its interaction with epistemic reasoning as well. The next section takes a first step in this direction: it adds a representation of the epistemic state of an agent to the model, extending the language with expressions that can talk about knowledge and knowledge-update in the context of causal reasoning.

⁸ The other two levels that Pearl distinguishes are the level of association, which is based on observation, and the level of intervention, which is based on doing. Modern AI technology is for him still at the first level: association. Counterfactual reasoning is not possible without a true understanding of *why* things happen – in our terminology, it is not possible without knowing the causal relationships as determined by \mathcal{F} .

There is another perspective from which such an epistemic extension of the standard framework can be motivated. In recent years there has been growing interest in the logic of dependence/determinacy. For instance, the IF logic of [22] expresses dependence by decorations of the quantifiers. Then, [30] and [2] use a primitive expression indicating that the value of one variable depends on that of another. In all these cases, the discussed notion of dependence/determinacy relies on considering a multiplicity of valuations in the model: the variable Y depends on (it is determined by) the variables X_1, \dots, X_n when, in all valuations that are being considered, fixing the value of the latter also fixes the value of the former. This gives rise to the question of how the notion of causal dependence modelled by the just introduced framework interacts with the notions of dependence/determinacy modelled by these alternative frameworks, and how causal dependence fits into a general picture of reasoning with and about dependencies. Interestingly, extending the standard causal reasoning approach with basic epistemic notions gives us another way to express the same notion of dependence as studied in the works just cited. This, then, allows us to compare different notions of dependency within one logical system. We will come back to this connection in Sect. 6.

4 Epistemic Causal Models

The first step towards a framework that combines causal with epistemic reasoning is adding a representation of the epistemic state of an agent to the causal model. This is done by adding a set of valuations \mathcal{T} , representing the alternatives the agent considers possible.

Definition 7 (Epistemic causal model). *An epistemic (note: recursive) causal model is a tuple $\langle \mathcal{S}, \mathcal{F}, \mathcal{T} \rangle$ where $\mathcal{S} = \langle \mathcal{U}, \mathcal{V}, \mathcal{R} \rangle$ is a signature, \mathcal{F} is a (note: recursive) set of structural functions for \mathcal{V} , and \mathcal{T} is a non-empty set of valuation functions for $\mathcal{U} \cup \mathcal{V}$, each one of them complying with \mathcal{F} .*

Example 1 can now be modelled as follows. We define an epistemic causal model $E = \langle \mathcal{S}, \mathcal{F}, \mathcal{T} \rangle$ whose signature \mathcal{S} has three variables: the exogenous B for the button and C for the circuit breaker, and the endogenous S for the sprinkler. All three variables can take two values, 0 or 1. The set of functions \mathcal{F} contains only one element: the function mapping S to 1 iff both B and C also have value 1. Because the agent can observe the value of the variables B and S , the set \mathcal{T} contains the assignment \mathcal{A}_1 that maps C to 0, B to 0 and S to 0, and the assignment \mathcal{A}_2 that maps C to 1, B to 0 and S to 0. Note how \mathcal{T} cannot contain the assignment $C = 1, B = 1$ and $S = 0$, for instance, because this assignment does not comply with the causal law in \mathcal{F} . This observation highlights an important feature of this notion of epistemic model: it cannot model uncertainty about the causal dependencies. Investigating the consequences of lifting this restriction is left for future research. The next step is to extend the notion of an intervention to epistemic causal models.

Definition 8 (Intervention). *Let $E = \langle \mathcal{S}, \mathcal{F}, \mathcal{T} \rangle$ be an epistemic causal model; let $\vec{X} = \vec{x}$ be an assignment on \mathcal{S} . The epistemic causal model $E_{\vec{X} = \vec{x}} = \langle \mathcal{S}, \mathcal{F}_{\vec{X} = \vec{x}}, \mathcal{T}_{\vec{X} = \vec{x}}^{\mathcal{F}} \rangle$, resulting from an intervention setting the values of variables in \vec{X} to \vec{x} , is such that*

- $\mathcal{F}_{\vec{X}=\vec{x}}$ is defined from \mathcal{F} just as in Definition 5,
- $\mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}} := \{\mathcal{A}'_{\vec{X}=\vec{x}}^{\mathcal{F}} \mid \mathcal{A}' \in \mathcal{T}\}$ (see Definition 5).

Note how $\langle \mathcal{S}, \mathcal{F}_{\vec{X}=\vec{x}}, \mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}} \rangle$ is indeed an epistemic causal model, as $\mathcal{F}_{\vec{X}=\vec{x}}$ is recursive and all valuations in $\mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}}$ comply with it.

In the just introduced model E for Example 1, we can now calculate the effects of considering the intervention that sets $B = 1$. According to Definition 8, an intervention on an epistemic causal model amounts to intervening on each of the assignments contained in the epistemic state. Thus, for our concrete example, we need to calculate the effects of an intervention with $B = 1$ on the assignments \mathcal{A}_1 and \mathcal{A}_2 that make up the epistemic state \mathcal{T} . The new epistemic state $\mathcal{T}_{B=1}^{\mathcal{F}}$ will now contain the assignment $\mathcal{A}_{1,B=1}^{\mathcal{F}}$ that maps C to 0, B to 1 and S to 0 and the assignment $\mathcal{A}_{2,B=1}^{\mathcal{F}}$ that maps C to 1, B to 1 and S to 1. Thus, the consequences of the intervention are calculated for all epistemic possibilities the agent considers. In other words, Definition 8 assumes that the agent has full epistemic access to the effect of the intervention on the model. In particular, she knows that the intervention takes place (in the counterfactual scenario considered). This makes a lot of sense if you think of the agent whose epistemic state is modelled as the one engaging in the counterfactual thinking. It is less plausible in connection to counterfactual thinking about the knowledge states of other agents. But this is something that we can leave for now, as we will not consider epistemic causal models for multiple agents in this paper. Based on these changes on the semantic side, we can now extend the object language with expressions that talk about the epistemic state of the agent. More specifically, we add the operator K for knowledge and “!” for information update. In other words, we understand “!” as expressing the action of observing or receiving information.

Definition 9. Formulas ϕ of the language \mathcal{L}_{PAKC} based on \mathcal{S} are given by

$$\begin{aligned} \gamma &::= Z = z \mid \neg\gamma \mid \gamma \wedge \gamma \mid K\gamma \mid [\gamma!]\gamma && \text{for } Z \in \mathcal{U} \cup \mathcal{V} \text{ and } z \in \mathcal{R}(Z) \\ \phi &::= Z = z \mid \neg\phi \mid \phi \wedge \phi \mid K\phi \mid [\phi!]\phi \mid [\vec{X}=\vec{x}]\gamma && \text{for } \vec{X}=\vec{x} \text{ an assignment on } \mathcal{S} \end{aligned}$$

Other Boolean operators (\vee , \rightarrow , \leftrightarrow) can be defined as usual. Note how, although the language makes free use of Boolean, epistemic and announcement operators (K and $[\phi!]$, for the latter two), nested intervention is again not allowed.⁹ Note also how the tuple vector \vec{X} can be empty, in which case $[\vec{X}=\vec{x}]\gamma$ becomes γ . The semantics for this extended language is straightforward.

Definition 10. Formulas of \mathcal{L}_{PAKC} are evaluated in a pairs (E, \mathcal{A}) with $E = \langle \mathcal{S}, \mathcal{F}, \mathcal{T} \rangle$ an epistemic causal model and $\mathcal{A} \in \mathcal{T}$. The semantic interpretation for Boolean operators is the usual; for the rest,

$$\begin{aligned} (E, \mathcal{A}) \models Z=z & \quad \text{iff} \quad \mathcal{A}(Z) = z \\ (E, \mathcal{A}) \models K\phi & \quad \text{iff} \quad (E, \mathcal{A}') \models \phi \text{ for every } \mathcal{A}' \in \mathcal{T} \\ (E, \mathcal{A}) \models [\psi!]\phi & \quad \text{iff} \quad (E, \mathcal{A}) \models \psi \text{ implies } (E^\psi, \mathcal{A}) \models \phi \\ (E, \mathcal{A}) \models [\vec{X}=\vec{x}]\gamma & \quad \text{iff} \quad (E_{\vec{X}=\vec{x}}, \mathcal{A}_{\vec{X}=\vec{x}}^{\mathcal{F}}) \models \gamma \end{aligned}$$

⁹ However, notice that the semantics already allows for nested occurrences of all dynamic operators. We will extend the proofs of sound- and completeness to the unrestricted language in the future.

with $E^\psi = \langle S, \mathcal{F}, \mathcal{T}^\psi \rangle$ such that $\mathcal{T}^\psi := \{\mathcal{A}' \in \mathcal{T} \mid (E, \mathcal{A}') \models \psi\}$. Note how E^ψ is an epistemic causal model: \mathcal{F} is recursive, and all valuations in \mathcal{T}^ψ comply with it.

Let us illustrate this definition with the help of the epistemic model E we introduced for Example 1. In order to evaluate a concrete formula with respect to this model we need to select, next to E , an assignment representing the actual world. In the example this is assignment \mathcal{A}_2 : in the actual world, the circuit breaker is closed, but because the button has not been pushed, the sprinkler is not working. We can calculate that the counterfactual $[B = 1]S = 1$ comes out as true given E and \mathcal{A}_2 , just as in the non-epistemic approach discussed in Sect. 2. But because we now also have a representation of the epistemic state of some agent, we can additionally consider epistemic attitudes the agent has towards this counterfactual. For instance, we can check that $K([B = 1]S = 1)$ is not true given E and \mathcal{A}_2 . For the sentence to be true, the formula $[B = 1]S = 1$ needs to be true over both (E, \mathcal{A}_1) and (E, \mathcal{A}_2) , because \mathcal{A}_1 and \mathcal{A}_2 are the two elements of \mathcal{T} . Thus, we need both $(E_{B=1}, \mathcal{A}_{1,B=1}^{\mathcal{F}}) \models S = 1$ and $(E_{B=1}, \mathcal{A}_{2,B=1}^{\mathcal{F}}) \models S = 1$. We already calculated $\mathcal{T}_{B=1}$ above: $\mathcal{T}_{B=1} = \{\mathcal{A}_{1,B=1}^{\mathcal{F}}, \mathcal{A}_{2,B=1}^{\mathcal{F}}\}$. But the content of $\mathcal{T}_{B=1}$ does not matter for the truth of the consequent $S = 1$ of the counterfactual that we are considering here, since this consequent does not contain epistemic operators. However, while in $\mathcal{A}_{1,B=1}^{\mathcal{F}}$ the sprinkler is still off, in $\mathcal{A}_{2,B=1}^{\mathcal{F}}$ it is on. This means that $(E_{B=1}, \mathcal{A}_{1,B=1}^{\mathcal{F}}) \not\models S = 1$, while $(E_{B=1}, \mathcal{A}_{2,B=1}^{\mathcal{F}}) \models S = 1$. Thus, the agent cannot predict the outcome of the intervention, just as intended in this case.

Finally, we define an operator \rightsquigarrow in terms of the existing vocabulary as a way to express causal dependency in the object language.

Definition 11. Take X and Z in $\mathcal{U} \cup \mathcal{V}$. The formula $X \rightsquigarrow Z$ is defined as

$$\bigvee_{\substack{\vec{w} \in \mathcal{R}((\mathcal{U} \cup \mathcal{V}) \setminus \{X, Z\}), \\ \{x_1, x_2\} \subseteq \mathcal{R}(X), x_1 \neq x_2, \\ \{z_1, z_2\} \subseteq \mathcal{R}(Z), z_1 \neq z_2}} [\vec{W}=\vec{w}, X=x_1]Z = z_1 \wedge [\vec{W}=\vec{w}, X=x_2]Z = z_2,$$

A formula $X \rightsquigarrow Z$ should be read as “ X has a direct causal effect on Z ”. It holds when there is a vector \vec{w} of values for variables in $\mathcal{R}(\mathcal{U} \cup \mathcal{V} \setminus \{X, V\})$ and two different values x_1, x_2 for X that produce two different values z_1, z_2 for Z (cf. [17]). When $Z \in \mathcal{V}$, it is clear that \rightsquigarrow is the syntactic counterpart of the relation “ \leftrightarrow ” of Definition 2.

5 Axiomatization

The axiom system L_{PAKC} is presented in Table 1. The *intervention* axioms, HP1-HP6, RH1 and RH2, are the standard axiomatization for the intervention operator over recursive causal models, with EX an additional axiom indicating that an exogenous variable is immune to interventions to any other variables. Then, the *epistemic* part contains the standard modal S5 axiomatization for truthful knowledge with positive and negative introspection.

Axiom CM indicates that what the agent will know after an intervention ($([\vec{X}=\vec{x}]K\phi)$) is exactly what she knows now about the effects of the intervention ($(K[\vec{X}=\vec{x}]\phi)$). Although maybe novel in the literature on causal models, the axiom is simply an instance of the more general DEL pattern of interaction between knowledge and a deterministic action without precondition. Axiom KR states that the agent knows how each endogenous variable $Y \in \mathcal{V}$ is affected when *all other variables* are intervened. Finally, axioms RP1-RP4 and rule RE in the *announcement* part are a *reduction-based* axiomatisation for public announcements in the DEL style. Here, axioms RP4 and RP1 are the most important. The first, RP4, is the well-known reduction axiom for announcement and knowledge, stating that knowing ϕ after an announcement of ψ is equivalent to knowing, conditionally on ψ , that the announcement of ψ would make ϕ true.¹⁰ The second, RP1, establishes the reduction for ‘atoms’ of the form $([\vec{X}=\vec{x}]Z = z)$; when \vec{X} is not empty, it states that a public announcement does not change the causal rules in the model.

Table 1. Axiom system \mathcal{L}_{PAKC}

<i>Propositional:</i>	
P $\vdash \phi$ for ϕ an instance of a tautology	MP From $\phi \rightarrow \psi$ and ϕ derive ψ
<i>Intervention:</i>	
HP1 $\vdash ([\vec{X}=\vec{x}]Z = z \rightarrow \neg[\vec{X}=\vec{x}']Z = z'$ for $z \neq z' \in \mathcal{R}(Z)$	
HP2 $\vdash \bigvee_{z \in \mathcal{R}(Z)} ([\vec{X}=\vec{x}]Z = z$	
HP3 $\vdash ([\vec{X}=\vec{x}]Z = z \wedge [\vec{X}=\vec{x}]W = w) \rightarrow ([\vec{X}=\vec{x}, Z=z]W = w$	
HP4 $\vdash ([\vec{X}=\vec{x}, Z=z]Z = z$	
HP5 $\vdash ([\vec{X}=\vec{x}, Z=z]W = w \wedge [\vec{X}=\vec{x}, W=w]Z = z) \rightarrow ([\vec{X}=\vec{x}]W = w$ for $W \neq Z$	
HP6 $\vdash (Z_0 \rightsquigarrow Z_1 \wedge \dots \wedge Z_{k-1} \rightsquigarrow Z_k) \rightarrow \neg(Z_k \rightsquigarrow Z_0)$	
RH1 $\vdash ([\vec{X}=\vec{x}](\gamma_1 \wedge \gamma_2) \leftrightarrow (([\vec{X}=\vec{x}]\gamma_1 \wedge [\vec{X}=\vec{x}]\gamma_2)$	
RH2 $\vdash ([\vec{X}=\vec{x}]\neg\gamma \leftrightarrow \neg[\vec{X}=\vec{x}]\gamma$	
EX $\vdash U = u \leftrightarrow ([\vec{X}=\vec{x}]U = u$ for $U \in \mathcal{U}$ with $U \notin \vec{X}$	
<i>Epistemic:</i>	
K $\vdash K(\phi \rightarrow \psi) \rightarrow (K\phi \rightarrow K\psi)$	T $\vdash K\phi \rightarrow \phi$
N From $\vdash \phi$ derive $\vdash K\phi$	4 $\vdash K\phi \rightarrow KK\phi$
	5 $\vdash \neg K\phi \rightarrow K\neg K\phi$
<i>Epistemic+Intervention:</i>	
CM $\vdash ([\vec{X}=\vec{x}]K\gamma \leftrightarrow K([\vec{X}=\vec{x}]\gamma)$	
KR $\vdash ([\vec{X}=\vec{x}]Y=y \rightarrow K([\vec{X}=\vec{x}]Y=y$ for $Y \in \mathcal{V}$ and $\vec{X} = (\mathcal{U} \cup \mathcal{V}) \setminus \{Y\}$	
<i>Announcement:</i>	
RP1 $\vdash [\psi!][\vec{X}=\vec{x}]Z = z \leftrightarrow (\psi \rightarrow ([\vec{X}=\vec{x}]Z = z)$	RP3 $\vdash [\psi!](\phi \wedge \chi) \leftrightarrow ([\psi!]\phi \wedge [\psi!]\chi)$
RP2 $\vdash [\psi!]\neg\phi \leftrightarrow (\psi \rightarrow \neg[\psi!]\phi)$	RP4 $\vdash [\psi!]K\phi \leftrightarrow (\psi \rightarrow K(\psi \rightarrow [\psi!]\phi))$
RE If $\vdash \psi_1 \leftrightarrow \psi_2$ and $\phi[\psi_2/\psi_1] \in \mathcal{L}_{PAKC}$, then $\vdash \phi \leftrightarrow \phi[\psi_2/\psi_1]$, where $\phi[\psi_2/\psi_1]$ is a formula obtained by replacing one or more non-announcement occurrences of ψ_1 in ϕ with ψ_2 . ^a	

A non-announcement occurrence of ψ in ϕ is an occurrence of ψ in ϕ where ψ is not inside the *brackets* of an announcement operator.

The axiom system \mathcal{L}_{PAKC} is sound and complete for \mathcal{L}_{PAKC} in epistemic causal models. Here is the argument for soundness.

Theorem 1. *The axiom system \mathcal{L}_{PAKC} is sound for \mathcal{L}_{PAKC} in epistemic causal models.*

¹⁰ Note how the announcement of ψ is a deterministic action *with precondition* ψ . Hence the similarities and differences between RP4 and CM.

Proof. For the soundness of HP1-HP6, RH1 and RH2 on causal models (enough for soundness on epistemic causal models, as evaluating the formulas does not require a change in valuation), see [17]. For the soundness of K, N, T, 4, and 5 on relational structures with an equivalence relation (equivalent to having a simple set of epistemic alternatives, as epistemic causal models have), see [10, 15]. For the soundness of RP1-RP4 when $[\psi!]$ describes the effect of a deterministic domain-reducing model operation, see [31].

For axioms EX, CM and KR, take any $(\langle S, \mathcal{F}, \mathcal{T} \rangle, \mathcal{A})$. For EX note how, for any $\vec{X} = \vec{x}$, the valuations \mathcal{A} and $\mathcal{A}_{\vec{X}=\vec{x}}^{\mathcal{F}}$ assign the same value to *exogenous* variables not occurring in \vec{X} (Definition 5). For CM, note how (i) $K[\vec{X}=\vec{x}]\phi$ holds at $(\langle S, \mathcal{F}, \mathcal{T} \rangle, \mathcal{A})$ iff ϕ holds at $(\langle S, \mathcal{F}_{\vec{X}=\vec{x}}, \mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}}, \mathcal{A}'_{\vec{X}=\vec{x}} \rangle, \mathcal{A}'_{\vec{X}=\vec{x}})$ for every $\mathcal{A}' \in \mathcal{T}$, and (ii) $[\vec{X}=\vec{x}]K\phi$ holds at $(\langle S, \mathcal{F}, \mathcal{T} \rangle, \mathcal{A})$ iff ϕ holds at $(\langle S, \mathcal{F}_{\vec{X}=\vec{x}}, \mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}}, (\mathcal{A}'_{\vec{X}=\vec{x}})' \rangle, (\mathcal{A}'_{\vec{X}=\vec{x}})')$ for every $(\mathcal{A}'_{\vec{X}=\vec{x}})' \in \mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}}$. Then it is enough to notice how, by Definition 8, the set of relevant valuations for the second, $\mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}}$, is exactly the set of relevant valuations for the first, $\{\mathcal{A}'_{\vec{X}=\vec{x}}^{\mathcal{F}} \mid \mathcal{A}' \in \mathcal{T}\}$. For KR, simply recall that all valuations in \mathcal{T} comply with the same structural functions. Finally, soundness of RE follows from two facts: the truth-value of every formula depends on the truth-value of its subformulas, and model operations (intervention and announcements) produce epistemic causal models. Thus, substituting a non-announcement subformula for a formula that is semantically equivalent in the given class of structures does not affect the final result. \square

The argument for completeness uses two steps. (i) First, using the reduction axioms technique, we show that L_{PAKC} allow us to translate any formula in \mathcal{L}_{PAKC} into a logically equivalent one without public announcements.¹¹ (ii) Then, relying on the canonical model construction for both causal models [17] and epistemic models [15, Chap. 3], we argue that L_{PAKC} is complete for the language without public announcements.

Theorem 2. *The axiom system L_{PAKC} is complete for \mathcal{L}_{PAKC} in epistemic causal models.*

6 Discussion

In this section we will compare our proposal to the Causal Team Semantics developed in [5–7]. Causal Team Semantics was proposed with the intention of supporting languages that discuss both accidental and causal dependencies. This is a topic that has gained quite some interest in recent years (see, e.g., [12, 21]). Causal Team Semantics was developed along the lines of a non-modal tradition of logics of dependence and independence (e.g. [22, 30]) by extending the so-called *team semantics* [20] with elements taken from causal inference. Even though the focus there is not on combining

¹¹ Readers familiar with DEL might have noticed that L_{PAKC} does not have a reduction axiom for nested announcements $[\phi_1!][\phi_2!]\phi$. There are (at least) two strategies for dealing with such formulas. The first follows an ‘outside-in’ approach, reducing two announcements in a row into a single one. This requires an axiom for nested announcements. The second follows an ‘inside-out’ strategy, applying the reduction over the innermost announcement operator in the formula until the operator disappears, and then proceeding to the next. For this, the rule of substitution of equivalents (our rule RE) is enough [31, Theorem 11].

causal with epistemic reasoning, this framework bears many similarities to the one we are using, which is why we will discuss it here in detail. Furthermore, this also allows us to say a bit more on the topic of dependence from the perspective of our proposal.

Let us quickly introduce the central notions of Causal Team Semantics to facilitate a comparison of the two frameworks. A causal team¹² is a tuple $T = \langle \mathcal{T}, \mathcal{F} \rangle$ where \mathcal{F} is defined similarly as in our paper¹³ and \mathcal{T} is a possibly empty set of valuations that comply with \mathcal{F} . Papers on Causal Team Semantics consider a variety of languages. The focus here is the one we shall call \mathcal{L}_{COD} , which is similar to the standard causal language (thus allowing to express various notions of causal dependence in terms of counterfactuals) except for the additional *dependence atoms* “ $=(X_1, \dots, X_n; Y)$ ”, which expresses (accidental) dependency of the variable Y on the variables X_1 to X_n . A sentence $=(X_1, \dots, X_n; Y)$ is interpreted as the claim that any two states s and s' that agree on the valuation of the variables X_1, \dots, X_n also have to agree on the value they assign to Y . Let us describe the syntax of \mathcal{L}_{COD} in more detail. The signatures used in [5] are pairs of the form $\langle Dom, Ran \rangle$, where Dom is a set of variables (*not* encoding the distinction between exogenous and endogenous variables) and Ran is defined analogously as the \mathcal{R} used in this paper. For any such fixed signature \mathcal{S} , the language \mathcal{L}_{COD} is defined as

$$\begin{aligned} \alpha &::= Z=z \mid Z\neq z \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \alpha \supset \alpha \mid \vec{X}=\vec{x} \mid \square \rightarrow \alpha \\ \phi &::= Z=z \mid Z\neq z \mid =(\vec{X}; Y) \mid \phi \wedge \phi \mid \phi \vee \phi \mid \alpha \supset \phi \mid \vec{X}=\vec{x} \mid \square \rightarrow \phi \end{aligned}$$

for $Z, Y, \vec{X} \in Dom$ and $z \in Ran(Z)$; and where the expression $\vec{X} = \vec{x}$ is an abbreviation for a conjunction of the form $X_1 = x_1 \wedge \dots \wedge X_n = x_n$.¹⁴ Below the complete semantics of \mathcal{L}_{COD} is given, using the notation of this manuscript. Notice that formulas are evaluated on a causal team *globally*: no valuation in \mathcal{T} is isolated as being ‘the actual world’. At the atomic level, this is done by means of a *universal* quantification. Indeed, while formulas of the form $Z = z$ and $Z \neq z$ indicate, semantically, that Z ’s value is (different from) z in *all valuations* in \mathcal{T} , a dependence atom $=(X_1, \dots, X_n; Y)$ indicates, as stated, that *all pairs* of valuations agreeing on the values of all X_i also agree on the value of Y . To keep the global perspective through the rest of the formulas, the interpretation of some connectives (\vee and \supset) differs from the traditional one (and, in particular, from that given on epistemic causal models). However, these connectives behave classically if applied to subformulas without occurrences of dependence atoms, and also when \mathcal{T} is a singleton (the quantification plays no relevant role).

¹² We are presenting here the definition from [7], which, save for implementation details, corresponds to what are called *fully defined* causal teams in [5] (where a more general notion is considered).

¹³ With some additional machinery, which is not worth exploring here.

¹⁴ Notice that the syntax allows negation only at the atomic level. Adding contradictory negation (defined by $T \models \sim\psi$ iff $T \not\models \psi$) would lead to a more expressive language and to an unintended reading of negation. As observed in [6], the language can be extended – without changes in expressivity – with a *dual negation*, defined by the clause: $(\mathcal{T}, \mathcal{F}) \models \neg\psi$ iff, for all $s \in \mathcal{T}$, $(s, \mathcal{F}) \not\models \psi$. The dual negation has the intended reading on formulas without dependence atoms. Neither negation allows the usual interdefinability of \wedge and \vee via the De Morgan laws; for this reason, both \wedge and \vee are included in the syntax.

$T \models Z=z$	iff	$s(Z) = z$ for all $s \in \mathcal{T}$
$T \models Z \neq z$	iff	$s(Z) \neq z$ for all $s \in \mathcal{T}$
$T \models =(X_1, \dots, X_n; Y)$	iff	for all $s, s' \in \mathcal{T}$, if $s(X_i) = s'(X_i)$ for $1 \leq i \leq n$, then $s(Y) = s'(Y)$
$T \models \phi \wedge \psi$	iff	$T \models \phi$ and $T \models \psi$
$T \models \phi \vee \psi$	iff	there are $\mathcal{T}_1 \cup \mathcal{T}_2 = \mathcal{T}$ such that $\langle \mathcal{T}_1, \mathcal{F} \rangle \models \phi$ and $\langle \mathcal{T}_2, \mathcal{F} \rangle \models \psi$
$T \models \alpha \supset \psi$	iff	$\langle \mathcal{T}^\alpha, \mathcal{F} \rangle \models \psi$, for $\mathcal{T}^\alpha := \{s \in \mathcal{T} \mid (\{s\}, \mathcal{F}) \models \alpha\}$ and α without dependence atoms
$T \models \vec{X}=\vec{x} \square \rightarrow \psi$	iff	$\langle \mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}}, \mathcal{F}_{\vec{X}=\vec{x}} \rangle \models \psi$, with $\mathcal{T}_{\vec{X}=\vec{x}}^{\mathcal{F}}$ and $\mathcal{F}_{\vec{X}=\vec{x}}$ as in Definition 8.

From their definitions, it is clear that an epistemic causal model and a causal team are identical objects; the only difference is that, for evaluating formulas, the former requires an ‘actual world’. On the syntactic side, even though the truth clauses of the logical operators differ in various respects, we can find several equivalences. For instance, the notion of dependence from team semantics can be expressed in our formal language as well.¹⁵ Indeed, interpret the object \mathcal{T} of a causal team as the epistemic state of some agent. Then, the statement $Y = y$ of causal team semantics can be understood as a claim about the knowledge of the agent, written in our language as $K(Y = y)$. Building on this translation, we can express that variable Y depends on the variables \vec{X} as the following claim: for all possible valuations \vec{x} of \vec{X} there is some value y of Y such that the agent knows that if she would observe $\vec{X} = \vec{x}$, she would know that Y has value y .

$$\bigwedge_{\vec{x} \in \mathcal{R}(\vec{X})} \bigvee_{y \in \mathcal{R}(Y)} [(X_1 = x_1 \wedge \dots \wedge X_n = x_n)!]K(Y = y).$$

With this idea in mind we can define a translation of the non-nested formulas of \mathcal{L}_{COD} .¹⁶ Setting aside for a moment the case of the operator \supset , and using A to denote the set of all possible valuations for $\mathcal{U} \cup \mathcal{V}$, the translation is given by the following clauses.

$$\begin{aligned} tr(Y=y) &:= K(Y=y) & tr(\phi_1 \wedge \phi_2) &:= tr(\phi_1) \wedge tr(\phi_2) \\ tr(Y \neq y) &:= K(\neg(Y=y)) & tr(\vec{X}=\vec{x} \square \rightarrow \phi) &:= [\vec{X}=\vec{x}]tr(\phi) \\ tr(\phi \vee \psi) &:= \bigvee_{S \subseteq A} K\left([\bigvee_{\vec{Y}=\vec{y} \in S} \vec{Y} = \vec{y}]!tr(\phi) \wedge [(\neg \bigvee_{\vec{Y}=\vec{y} \in S} \vec{Y} = \vec{y})!]tr(\psi)\right) \\ tr(=(X_1, \dots, X_n; Y)) &:= \bigwedge_{\vec{x} \in \mathcal{R}(\vec{X})} \bigvee_{y \in \mathcal{R}(Y)} [(X_1=x_1 \wedge \dots \wedge X_n=x_n)!]K(Y=y) \end{aligned}$$

A short note on the not-so-intuitive translation clause for \vee . First note that, in the semantic clause for \vee , the sets \mathcal{T}_1 and \mathcal{T}_2 can equivalently be required to form a *partition*

¹⁵ As far as we know, this has been first observed, independently, in [14] and [1], in the context of epistemic languages with modalities for the knowledge of values.

¹⁶ A formula is non-nested if, in every subformula of the form $\vec{X}=\vec{x} \square \rightarrow \phi$, no $\square \rightarrow$ occurs inside ϕ . Providing a translation for these formulas is sufficient, since every formula of the causal team language is provably equivalent to a non-nested one.

of \mathcal{T} , i.e. to be disjoint. The translation clause uses the fact that a partition of \mathcal{T} (say, $\mathcal{T} \cap S$ and $\mathcal{T} \setminus S$) can be characterized by the pair of formulas $\bigvee_{\vec{y}=\vec{y} \in S} \vec{Y} = \vec{y}$ (defining $\mathcal{T} \cap S$ as a subset of \mathcal{T}) and $\neg \bigvee_{\vec{y}=\vec{y} \in S} \vec{Y} = \vec{y}$ (defining $\mathcal{T} \setminus S$). The conjunction $[(\bigvee_{\vec{y}=\vec{y} \in S} \vec{Y} = \vec{y})!]tr(\phi) \wedge [(\neg \bigvee_{\vec{y}=\vec{y} \in S} \vec{Y} = \vec{y})!]tr(\psi)$ then ensures that the current assignment either is in S and satisfies the translation of ϕ , or it is in $\mathcal{T} \setminus S$ and satisfies the translation of ψ . The K operator, placed *after* the disjunction $\bigvee_{S \subseteq A}$, ensures that, fixing a partition, this property holds for all the assignments (i.e. the partition is not picked out as a function of the assignment). Notice also that this translation clause – as well as that for dependence atoms – is well-defined relative to a fixed, finite signature, since the translation uses an enumeration of the variables and of their corresponding allowed values.

Formulas of the form $\alpha \supset \psi$ translate into public announcement formulas. However, in order to play the role of announcement, α cannot be translated using tr , as announcements are evaluated according to the classical meaning. We need instead a simpler translation e which just replaces logical operators with their counterparts in \mathcal{L}_{PAKC} ($X \neq x$ is replaced by $\neg(X = x)$; $\beta \supset \gamma$ by $\beta \rightarrow \gamma$; $\vec{X} = \vec{x} \square \rightarrow \phi$ by $[\vec{X} = \vec{x}]\phi$; \wedge and \vee are left unaltered, or, more precisely, $\beta \vee \gamma$ is replaced by $\neg(\neg\beta \wedge \neg\gamma)$). Then we can define tr for \supset as follows:

$$tr(\alpha \supset \phi) := [e(\alpha)!]tr(\phi)$$

This translation satisfies the following.

Proposition 1 (Global translation). *For any causal team $\langle \mathcal{T}, \mathcal{F} \rangle$ over a finite signature S and any formula $\phi \in \mathcal{L}_{COD}$, we have $\langle \mathcal{T}, \mathcal{F} \rangle \models \phi$ if and only if, for all $\mathcal{A} \in \mathcal{T}$, we have $(\langle S, \mathcal{F}, \mathcal{T} \rangle, \mathcal{A}) \models tr(\phi)$.*

This result compares truth on a causal team with validity over an epistemic causal model. On the other hand, a different translation of the dependence atom from [1, 14] suggests an alternative, “local” translation. Let tr^* be as tr , except for the following clauses (notice the additional K operator in both clauses):

$$tr^*(=(X_1, \dots, X_n; Y)) := \bigwedge_{\vec{x} \in \mathcal{R}(\vec{X})} \bigvee_{y \in \mathcal{R}(Y)} K[(X_1 = x_1 \wedge \dots \wedge X_n = x_n)!]K(Y = y)$$

$$tr^*(\alpha \supset \phi) := K[e(\alpha)!]tr(\phi)$$

Now we have the following result.

Proposition 2 (Local translation). *For any causal team $\langle \mathcal{T}, \mathcal{F} \rangle$ over a finite signature S and any formula $\phi \in \mathcal{L}_{COD}$, we have:*

- (i) *If $\langle \mathcal{T}, \mathcal{F} \rangle \models \phi$, then, for all $\mathcal{A} \in \mathcal{T}$, $(\langle S, \mathcal{F}, \mathcal{T} \rangle, \mathcal{A}) \models tr(\phi)$.*
- (ii) *If there is an $\mathcal{A} \in \mathcal{T}$ such that $(\langle S, \mathcal{F}, \mathcal{T} \rangle, \mathcal{A}) \models tr(\phi)$, then $\langle \mathcal{T}, \mathcal{F} \rangle \models \phi$.*

This result shows that, *in the finite case*, \mathcal{L}_{PAKC} is at least as expressive as \mathcal{L}_{COD} . Despite this, the way the notion of (accidental) dependence is spelled out in the two languages differs in an interesting way. While it is a primitive element in the language of Causal Team Semantics, the way it is definable in our epistemic framework emphasises what we can *do* with such a concept of dependence: we can make predictions based on what we observe. Furthermore, it is interesting to notice the similarity between this translation of (accidental) dependence and the way causal dependence is expressed. It is also not defined as a primitive in the language, but can be expressed using counterfactuals, which work based on the concept of intervention. These counterfactuals, in turn, focus on what you can do with causal information: prediction based on intervention.

Based on the counterfactual expression, various notions of causal dependence can be defined. We saw one already in Sect. 4, Definition 11: $X \rightsquigarrow Z$, which expresses that X is a causal parent of Z (if Z is an endogenous variable). The local translation of the notion of dependence from Causal Team Semantics into our framework suggests a different notion of causal dependence. We repeat the local translation below under the name of e-dependence. C-dependence defines the corresponding causal notion.¹⁷

- Y *e*-depends on X in (E, \mathcal{A}) iff $(E, \mathcal{A}) \models \bigwedge_{x \in \mathcal{R}(X)} \bigvee_{y \in \mathcal{R}(Y)} K([(X = x)!]K(Y = y))$
- Y *c*-depends on Y in (E, \mathcal{A}) iff $(E, \mathcal{A}) \models \bigwedge_{x \in \mathcal{R}(X)} \bigvee_{y \in \mathcal{R}(Y)} [X = x]K(Y = y)$

Given an epistemic causal model, C-dependence holds between a list of variables X_1, \dots, X_n and a variable Y if any intervention fixing the value of the variables X_1, \dots, X_n also determines the value of Y *within the epistemic state of the agent*. While this notion is certainly more robust than the notion of e-dependence, it still takes into account the epistemic state of the agent. The less the agent knows about the values of the variables, the more variables she needs to control to make sure that a variable Y is in a particular state. If the agent knows more about the actual causal history of Y , she can predict the state of Y already from smaller interventions. These kind of hybrid notions between causal and epistemic dependence that our framework allows to define deserve certainly some attention in future research.

7 Conclusions

In this paper we have moved some steps towards the integration of causal and epistemic reasoning, providing an adequate semantics, a language combining interventionist counterfactuals with (dynamic) epistemic operators and a sound and complete system of inference. Our deductive system models the thought of an agent reasoning about the consequences of hypothetical interventions and observations. It describes what the agent may deduce from her/his *a priori* pool of knowledge about a system of variables. It is therefore a logic of thought experiments. Going back to Example 1 from the introduction, the approach allows us to account for the inference that Billie is not sure that

¹⁷ The additional K operator in the definition of e-dependence is needed to deal with the fact that information update always checks first whether the information that the information state is updated with is true. This problem disappears in the case of interventions, because the formula you intervene with is *made* true in the hypothetical scenario you consider.

if the button had been pushed, the sprinkler would have been working. However, the logic is not yet able to also model the second inference discussed in connection with this example: if Billie had pushed the button and saw that the sprinkler works, then she would have known that the circuit is closed. In order to account for this kind of reasoning we need to model how an agent may reason about (from her perspective) actual experiments. Things change significantly in such a setting: because of unobserved factors, the agent may fail to predict the outcome of an experiment; yet the outcome may sometimes be recovered from direct observation of the consequences of the experiment. The development of a such a framework will involve a more careful distinction between *observable* and *unobservable* variables. The resulting logic must necessarily abandon the right-to-left implication of axiom CM ($([\vec{X}=\vec{x}]K\phi \rightarrow K[\vec{X}=\vec{x}]\phi)$, which expresses the fact that interventions cannot increase the knowledge of the agent.

Our framework has many points in common with the earlier causal team semantics, and we provided a translation between the two approaches. For the purpose of modeling causal reasoning, our semantics has the advantage, over causal team semantics, of encoding explicitly a notion of actual state of the world (and in particular, of actual value of variables). Actual values seem to be crucial for the attempt of defining notions of *token causation* [18, 19, 32], i.e. causation between events. In order to fully appreciate this advantage, though, we will need to consider richer languages with hybrid features that allow to explicitly refer to the actual values of variables.

Finally, in future work we plan to extend the setting to a multi-agent system. This involves considering not only different agents with potentially different knowledge, but also epistemic attitudes for groups (e.g., distributed and common knowledge) and the effect of inter-agent communication. One advantage this will bring is the potential to contribute to the discussion about causal agency and the role of causation in the study of responsibility within AI (f.i. [3]).

References

1. Baltag, A.: To know is to know the value of a variable. *Adv. Modal Log.* **11**, 135–155 (2016)
2. Baltag, A., van Benthem, J.: A simple logic of functional dependence. *J. Philos. Log.* (2020)
3. Baltag, A., Canavotto, I., Smets, S.: Causal agency and responsibility: a refinement of STIT logic. In: Giordani, A., Malinowski, J. (eds.) *Logic in High Definition. Trends in Logical Semantics* (forthcoming)
4. Baltag, A., Moss, L.S., Solecki, S.: The logic of public announcements, common knowledge, and private suspicions. In: Gilboa, I. (ed.) *TARK*, pp. 43–56. Morgan Kaufmann, San Francisco (1998). <http://dl.acm.org/citation.cfm?id=645876.671885>
5. Barbero, F., Sandu, G.: Interventionist counterfactuals on causal teams. In: Finkbeiner, B., Kleinberg, S. (eds.) *Proceedings 3rd Workshop on Formal Reasoning about Causation, Responsibility, and Explanations in Science and Technology*, Thessaloniki, Greece, 21st April 2018. *Electronic Proceedings in Theoretical Computer Science*, vol. 286, pp. 16–30. Open Publishing Association (2019)
6. Barbero, F., Sandu, G.: Team semantics for interventionist counterfactuals: observations vs. interventions. *J. Philos. Log.* (2020, to appear)
7. Barbero, F., Yang, F.: Counterfactuals and dependencies on causal teams: expressive power and deduction systems. *Adv. Modal Log.* (2020, accepted, forthcoming)

8. van Benthem, J.: *Logical Dynamics of Information and Interaction*. Cambridge University Press, Cambridge (2011)
9. Bergstein, B.: What AI still can't do (2020). <https://www.technologyreview.com/2020/02/19/868178/what-ai-still-cant-do/>. Accessed June 2020
10. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. No. 53 in *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge (2001). <https://doi.org/10.1017/CBO9781107050884>
11. Briggs, R.: Interventionist counterfactuals. *Philos. Stud.* **160**(1), 139–166 (2012)
12. Chockler, H., Halpern, J.Y.: Responsibility and blame: a structural-model approach. *J. Artif. Int. Res.* **22**(1), 93–115 (2004)
13. van Ditmarsch, H., van der Hoek, W., Kooi, B.: *Dynamic Epistemic Logic*. Synthese Library Series, vol. 337. Springer, Dordrecht (2008). <https://doi.org/10.1007/978-1-4020-5839-4>
14. van Eijck, J., Gattinger, M., Wang, Y.: Knowing values and public inspection. In: Ghosh, S., Prasad, S. (eds.) *ICLA 2017*. LNCS, vol. 10119, pp. 77–90. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54069-5_7
15. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: *Reasoning About Knowledge*. The MIT Press, Cambridge (1995)
16. Galles, D., Pearl, J.: An axiomatic characterisation of causal counterfactuals. *Found. Sci.* **1**, 151–182 (1998)
17. Halpern, J.Y.: Axiomatizing causal reasoning. *J. Artif. Intell. Res.* **12**, 317–337 (2000)
18. Halpern, J.Y.: *Actual Causality*. MIT Press, Cambridge (2016)
19. Hitchcock, C.: A tale of two effects. *Philos. Rev.* **110**(3), 361–396 (2001). <https://doi.org/10.1215/00318108-110-3-361>
20. Hodges, W.: Compositional semantics for a language of imperfect information. *Log. J. IGPL* **5**, 539–563 (1997)
21. Ibeling, D., Icard, T.: Probabilistic reasoning across the causal hierarchy. arXiv preprint [arXiv:2001.02889](https://arxiv.org/abs/2001.02889) (2020)
22. Mann, A.L., Sandu, G., Sevenster, M.: *Independence-Friendly Logic - A Game-Theoretic Approach*. London Mathematical Society Lecture Note Series, vol. 386. Cambridge University Press, Cambridge (2011)
23. Pearl, J.: *Causality. Models, Reasoning, and Inference*. Cambridge University Press, Cambridge (2000)
24. Pearl, J.: Causal diagrams for empirical research. *Biometrika* **82**(4), 669–688 (1995)
25. Pearl, J.: Causality: models, reasoning, and inference. *IIE Trans.* **34**(6), 583–589 (2002)
26. Pearl, J.: *Causality. Models, Reasoning, and Inference*, 2nd edn. Cambridge University Press, Cambridge (2009)
27. Pearl, J., Mackenzie, D.: *The Book of Why. The New Science of Cause and Effect*. Penguin Books, London (2019)
28. Spirtes, P., Glymour, C., Scheines, R.: *Causation, Prediction, and Search*, 1st edn. MIT Press, Cambridge (1993)
29. Strotz, R.H., Wold, H.O.: Recursive vs. nonrecursive systems: an attempt at synthesis (part I of a triptych on causal chain systems). *Econ.: J. Econ. Soc.* 417–427 (1960)
30. Väänänen, J.: *Dependence Logic: A New Approach to Independence Friendly Logic*. London Mathematical Society Student Texts, vol. 70. Cambridge University Press, Cambridge (2007)
31. Wang, Y., Cao, Q.: On axiomatizations of public announcement logic. *Synthese* **190**(1), 103–134 (2013)
32. Woodward, J.: *Making Things Happen*. Oxford Studies in the Philosophy of Science, vol. 114. Oxford University Press, Oxford (2003)