Modelling linguistic causation

Rebekah Baglini and Elitzur A. Bar-Asher Siegal*

March 2025

Abstract

This paper develops a formal methodology for capturing and representing the semantics of causal expressions in natural languages. Focusing on two causative constructions—covert causatives (change-of-state verbs) and overt causatives (the verb *cause*)—it provides a proof of concept for analyzing the distinguished meanings of different causative constructions.

We adopt the formal framework of Structural Equation Modeling (SEM) to analyze causality and integrate it into model-theoretic semantics for interpreting causal statements. In our approach, the selection of a cause within a particular construction depends on its inclusion in *a sufficient set of conditions* that bring about the effect, as well as on specific properties of the cause itself. To formalize this process, we introduce the concept of causative-construction selection (CC-selection), which captures how speakers select a causative construction that aligns with the relational structure between states of affairs. For each relevant condition within the sufficient set, CC-selection determines whether it can be encoded as the cause in a statement articulated through a specific causative construction, thereby describing a particular state of affairs. We argue that CC-selection plays a central role in shaping the meaning of causal statements.

By leveraging the SEM framework, CC-selection effects can be formally explained through contrasts within the structure of a model. For instance, notions of *sufficiency* and *necessity*, which play a crucial role in these selections, are rigorously defined within SEM, allowing for a precise account of CC-selection effects. This paper further illustrates how CC-selection accounts for contrastive inference patterns across constructions. By focusing on the two causative constructions central to our discussion, it resolves longstanding puzzles associated with change-of-state verbs. The proposed framework establishes a foundation for the systematic study of causal language, bridging semantics and philosophy while providing tools to investigate the interplay between causative constructions and their associated causal meanings.

1 Introduction

The occurrence of any event depends on various factors. More precisely, a target event happens as a result of a combination of conditions, each of which may be individually necessary but only collectively sufficient to produce the desired outcome. To illustrate, consider a door that can be opened in two ways: (a) through an automated process requiring electricity, an unlocked door, and someone pressing the door-open button; or (b) through a manual process involving a door handle, an unlocked door, and someone turning the handle. Each of these routes comprises a set

*The material presented in this paper was taught and discussed in seminars at several academic institutions, including Aarhus University, the Hebrew University of Jerusalem, the University of Chicago, Yale University, and the Summer Institute of the Linguistic Society of America. We are grateful for the insightful comments and suggestions provided by the participants in these seminars. We also sincerely thank the anonymous reviewers for their thoughtful feedback, which has significantly improved this paper. Research on this paper was supported by the Volkswagen Stiftung, "Forschungskooperation Niedersachsen-Israel" for the project "Talking about causation: linguistic and psychological perspectives" given to Elitzur Bar-Asher Siegal (together with Nora Boneh and York Hagmayer) as well as a Starting Grant from Aarhus Universitets Forskningsfond awarded to Rebekah Baglini.

of conditions that together are sufficient to open the door, and within each set, every mentioned condition is necessary. Now, imagine a situation where a person approaches the door, pushes the button, and the door opens. An observer attempting to describe this event in causal terms, i.e., to establish the dependency between different states of affairs in the world, faces several challenges. Natural language provides various causative constructions (e.g. verbs, connectives; see Table 1 below) to describe such relationships. However, linguistically, these constructions usually represent a binary relationship, where a single cause is linked to an effect (in this case, the opening of the door).[1] Thus, the observer encounters the problem of **causal selection**: deciding which of the known necessary and jointly sufficient conditions should be identified as the primary cause of the door's opening.

Causal selection has been extensively discussed in philosophy and the cognitive sciences, and we will delve further into it in Section 2.1. Existing research often treats causal selection as a monolithic cognitive operation with consistent linguistic outcomes. This perspective suggests a universal mechanism of causal identification, positing that the same cognitive process is invoked across various languages and linguistic structures to identify the same cause in any given scenario. This line of reasoning can be associated with a broader assumption found in much of the philosophical literature, which traces back to David Hume, that there should be a **unified account** for attributing the terms "cause" and "effect" to two entities (events or individuals).[2]

However, less attention has been given to a second, equally important selection problem faced by our hypothetical observer, recently referred to as **causative-construction selection (CC-selection)** (Bar-Asher Siegal et al., 2021). In generating a linguistic description of the situation, the observer needs not only to designate a particular (necessary) condition as the cause, but also must determine which linguistic construction appropriately describes the relationship underlying the observed sequence of events. Assuming that, in the previous task of causal selection, the observer selects pushing the button as the cause, two highly plausible alternatives to describe the event arise:

(1) Pushing the button opened the door.

(2) Pushing the button caused the door to open.

Linguistic studies often assume that the category of **causative constructions**[3] (exemplified in Table 1) is defined by their semantic property of denoting causal relations (Shibatani, 1976; Dowty, 1979; Comrie, 1981; Escamilla Jr, 2012). The theoretical implication of accepting both **the unified philosophical approach to causation** (that there is only one way in which two events/individuals can be causally dependent on each other) and the linguistic assumption about **the denotation of causative expressions** (that causative constructions denote causal relations) would be as follows:

[1] As will become clear, according to our analysis, the need for a binary cause-effect relationship is only imposed by the constraints of linguistic expressions, and is not something that is true about the nature of cognitively-perceived causal relationships themselves.

[2] It is not within the scope of the current paper to present the extensive history of discussions on causation in the philosophical literature. Bar-Asher Siegal and Boneh (2020) (especially on pp.1-12) provide a recent review of several significant topics discussed in this literature, including comments on their relevance to linguistic discussions on causation.

[3] Following Bar-Asher Siegal and Boneh (2020), by causative constructions we mean a semantically distinguished set of linguistic forms (including but not limited to those in Table 1) which encode a dependency between causes and effects with the following three components:

i) a cause (c);
ii) the effect of the cause (e); and
iii) the dependency (D) between c and e:
[c] D [e]

The terms "cause" (c) and "effect" (e) are used here loosely in a pre-theoretical manner. The use of the term "causative" or the division of the components to "cause" and "effect", at this point, neither indicates an assumption that a construction denotes causal relations, nor does it commit to the nature of (c) and (e). (c), (e) and D are used here in an uncommitted manner, relying on a pre-theoretical intuitions that they express some causal dependency, and it is our goal to understand the nature of these dependencies.

Causative constructions	
Dedicated verbs:	*cause, make, allow, enable, get...*
Connectives:	*because (of), from, by, as a result of..*
Change-of-state verbs:	*open, boil...*
Dedicated Morphology:	*C-templates* (Semitics), *suffix -ita* (Korean)

Table 1: Causative constructions

> If all causative constructions denote causal relations, and if all causal relations lend themselves to a unified account, then **the semantics of all causative constructions (with the same cause and the same effect) should be the same.**

This way of reasoning is reflected in the common assumption within the linguistic literature that there exists only a single type of dependency to be uncovered (such as Dowty's (1979) primitive predicate CAUSE, for instance), and it is often reflected in studies in semantics and syntax, when they assume that causative constructions share causal components.[4] Indeed, philosophers and cognitive psychologists who analyse causation based on causal judgements rarely factor in potential differences between the causative constructions in their research.

However, linguistic research has shown that different causative constructions have distinct meanings (for various inferential differences between constructions, see Shibatani (1976); Dixon (2000); Maienborn and Herdtfelder (2017); Nadathur and Lauer (2020); Bar-Asher Siegal and Boneh (2019, 2020)). For example, while a change-of-state (=CoS) causative (such as transitive 'open' in (1)) entails the truth of an overt *cause* sentence, an entailment in the other direction does not hold.

(3) a. Sam opened the door. $\models$ Sam caused the door to open.
b. Sam caused the door to open. $\not\models$ Sam opened the door.

The asymmetrical entailment patterns in (3) illustrate that, despite their apparent similarity in meaning, a CoS causative such as "open" and its periphrastic alternative with *cause* are not semantically equivalent (Hall, 1965, 28). Intuitively, this asymmetry is often attributed to a requirement of "directness" in CoS causatives. For instance, sentence (3-b)—but notably not (3-a)—can describe a scenario in which Sam opened a window, causing a gust of wind that subsequently pushed the door open.[5] Thus, the linguistic perspective on causal language leads to a conclusion that contradicts the earlier assumption:

> **It is not the case that the semantics of all causative constructions (with the same cause and the same effect) is the same.**

This apparent contradiction between the assumptions in the literature can be resolved by either discarding or revising one of them. One alternative is to abandon the idea of a unified account for causation, which assumes that causation is a singular type of relation or connection between things in the world. Instead, one can advocate **a pluralistic notion of causation**, which involves recognizing that the seemingly simple and unambiguous term "cause" conceals underlying diversity (for such approaches in philosophy, see Hitchcock (2003); Hall (2004); and in cognitive psychology: Waldmann and Hagmayer (2013)) From a linguistic perspective, taking this approach entails providing an independent analysis for each structure, without attempting to present a unifying element that encompasses a comprehensive analysis of causality in general (this direction was proposed by Copley and Wolff (2014), Bar-Asher Siegal and Boneh (2019) and Rose et al. (2021)).[6]

[4]For examples from the linguistic literature for this tendency to assume a shared causal components across constructions see Bar-Asher Siegal and Boneh, 2020, 10-11.

[5]The acceptability judgments and the conditions under which various sentences in the article are considered acceptable or unacceptable have been extensively tested with English native speakers in numerous presentations, who have verified these assessments.

[6]There are several reasons to favor a unified approach. Firstly, from an aesthetic perspective, adhering to the principle

The second approach involves refining the assumption about the relationship between causative constructions and causal relations. This approach seeks to integrate all relevant causal relations underlying the semantics of causal language into a unified concept of causation, while simultaneously introducing construction-specific constraints. Examples of this strategy can be found in works such as Shibatani (1976) and Levshina (2016), among others. However, as will be shown, these studies face challenges in capturing the inferential differences between constructions within a model-theoretic framework. Specifically, they lack a formal mechanism to represent the nuanced relationships between dependencies.[7]

This paper takes the latter approach and adopts a systematic method for elucidating the distinctions among the semantics of various causative constructions, exploring how variations in the meanings of linguistic expressions can be connected to causal models that encapsulate human knowledge of causality. Unlike approaches that treat causality as a uniform binary relation between a cause and an effect, we argue that causality is better understood as a complex system of conditions that collectively bring about an outcome. This complexity allows for the semantic diversity observed across different causative constructions, each of which imposes unique constraints tied to distinct aspects of this intricate system.

To this end, we adopt a formal framework to encapsulate the semantics of causative linguistic expressions, clarifying how causal relationships are represented in natural languages. One of the primary contribution of this work is the establishment of a framework wherein each causative linguistic construction is governed by unique constraints. These constraints specify which conditions, derived from a causal model, are recognized as the cause in any given causal statement. **We propose that these constraints serve as the truth conditions for using such linguistic expressions, offering a robust method for capturing their semantics.** This development involves two primary tasks:

> **Task 1:** Modelling causality and developing a corresponding semantic framework to capture the meaning of causative constructions within this model (this will unravel the relationship between the causative expressions and causal relations.)
> **Task 2:** Delineating the differences between causative constructions in a systematic and principled way, and and examining how these differences can be captured within the formalized models (Task 1).

To achieve these objectives, we adopt the formal framework of Structural Equation Modelling (SEM) adapted to represent causal relations (for our purposes we group SEM and causal Bayesian networks (CBN) together). This approach has been influential across a wide range of fields, including computer science, statistics, engineering, epidemiology, and philosophy and psychology (Pearl, 2000; Spirtes et al., 2000; Steyvers et al., 2003). Such studies use directed acyclic graphs to model causality, by representing dependencies between states of affairs as dependencies between valued variables. These models treat various qualitative notions of causal dependence, which philosophers such as Woodward (2003) and Hitchcock (2020), have used to account for the meaning or content of causal primitives. In the past decade, several works in linguistics have also proposed the use of the SEM approach to model the truth conditions of various linguistic expressions. Schulz (2011), Henderson (2014), Snider and Bjorndahl (2015) and Ciardelli et al. (2018), employ causal models for analyzing counterfactuals. Baglini and Francez (2016) and Nadathur and Lauer (2020) utilize these models for lexical semantics. Our approach largely aligns with the latter two papers.

This paper sets out with ambitious goals, aiming to develop a methodology for capturing and representing the semantics of causal expressions in language. While our discussion will be pri-

of Occam's razor, suggests that theories positing fewer entities or types of dependencies, are preferable to those positing more. In our case, there is a more compelling reason supporting a unified approach. It can be observed that causative constructions are often interchangeable, and different constructions can be used to describe the same state of affairs. This observation implies that these constructions share a fundamental aspect in their core meaning.

[7] A similar analytical approach is evident in Dixon (2000), who conceptualizes causation as a relationship between individuals, wherein one individual initiates or controls the activity. Dixon (2000) also argues that when a language exhibits multiple constructions, there must be a semantic distinction between them. According to his analysis, these distinctions may arise from either the verb itself, specifically the type of eventuality it denotes, or from the participants involved in the relationship, such as their manners of participation in the eventuality.

marily focused on two causative constructions, this study should be viewed as a proof of concept, illustrating how to formally distinguish between meanings across causative constructions. Thus, our actual analysis will concentrate on exploring the asymmetry among the causative constructions introduced in (3), a topic that is notably contentious within the linguistic literature on causal language. Our objective is to formally represent how the meanings of the two constructions interact with causal models and to illustrate this representation using the SEM framework. We will detail the distinct truth conditions associated with each construction, thereby elucidating their inferential relationships.

Given that our formal analysis is grounded in the SEM framework, it is necessary to introduce this framework at the outset. Accordingly, the exposition begins with two sections dedicated to accomplishing Task 1. Section 2 addresses the problem of causal selection and introduces the SEM framework from this perspsctive. This foundation is crucial for demonstrating our central claim that the semantics of causal expressions are shaped by constraints on the selection of specific conditions within a causal model. Section 3 provides the formal definitions that underpin our analysis. While building on an existing framework, this section offers a novel perspective on which formal aspects of causal models are relevant to linguistic expressions and introduces new formal definitions tailored to these elements.

Moving forward, Section 4 focuses on the two constructions that form the cornerstone of our proof of concept. We engage with the debate in the literature regarding the semantic asymmetry between overt and CoS causative constructions, examining how they represent causal relationships as either direct or indirect and addressing the lack of a comprehensive empirical solution. Section 5 employs formal semantic analysis to unravel this complexity, offering a case study that illustrates the practical application of Task 2. This analysis uses the formal definitions which are developed in Section 3. This method is consistent with related research (referenced throughout the paper,) and we foresee its extension to the analysis of other causative expressions. The paper culminates in Section 6, where we encapsulate our findings and contributions, providing a summary of our study.

2 Modelling Causation

2.1 Causal Selection and Causative Construction Selection

When a speaker observes that Sam pressed a button, which was followed by a door opening, and subsequently states, "Sam opened the door," the statement typically does not aim to convey information about the causal mechanism. Instead, it focuses on the agent, the action, and the resulting outcome. In everyday discourse, speakers and listeners generally share a common understanding of the causal mechanisms underlying such events, which informs their communication. Consequently, the utterance "Sam opened the door" is more likely intended to emphasize the agent's role or the action performed, rather than to explicate the causal relationship between pressing the button and the door's movement. This shared understanding implies a shared causal model of how the world operates.[8] According to this perspective, causal statements rely on causal knowledge and may not necessarily be directly meant to inform **about** the causal relations in the world (i.e. causality is not *at-issue*).[9] We argue that, in all cases—whether they establish causality or depend on causal knowledge—causal statements are true **in virtue of** the causal model that represents the causal knowledge of the interlocutors.

Therefore, we aim to develop a model of causality that enhances our understanding of the semantics embedded in causative constructions. This involves delving into the components of these

[8]Our focus here is on formalizing the already established causal model that speakers have as the background for their conversations, rather than describing the dynamics of its construction. This differs from aims like those of Pearl (2000), who seeks to delineate the structure that scientists aim to identify from real-world observations. The differing perspectives on referencing the model are discussed further in Section 5.1.

[9]See Bar-Asher Siegal and Boneh (2020) for a discussion on whether a causative construction explicitly asserts or merely implies a causal relation, and for the possibility that not all causative constructions behave uniformly in this regard—some explicitly assert causality, while in others, it is merely assumed.

models, scrutinizing their formal structures, and presenting an analysis that links these models to the meanings conveyed by causal expressions. The objective of this section is to accomplish these goals, thereby initiating the first task of this paper. Before proceeding, it is important to briefly outline how our approach differs from previous treatments of the semantics of causative constructions.

At the heart of the discourse on causation is the premise that causation is a binary relationship between a singular cause and its subsequent effect. This assumption also aligns with a widely accepted philosophical standpoint,[10] and it underpins existing semantic analyses of causal expressions, which have predominantly focused on delineating the relation between the cause and the effect. These analyses frequently use the primitive predicate CAUSE to represent causal interactions, as illustrated in Lakoff (1970), Jackendoff (1972), Levin and Rappaport-Hovav (1995) and Pylkkänen (2008).

Contrary to this entrenched assumption in the literature, we propose the hypothesis that the perceived limitation of binary relations—comprising a singular cause and effect—is more aptly ascribed to the linguistic forms themselves, which enforce this dichotomy, as they describe dependencies between events, and they are constructed as such (cf. Hitchcock 2020). However, causality itself is not merely a binary relationship. Philosophers like John Stuart Mill have illuminated that real-world causality is a complex notion involving the interrelations among multiple factors and their collective outcomes. Causal statements invariably emerge from a process of *causal selection*—the act of distinguishing one singular factor as "The Cause" amid a multitude of contributory factors. For example, in a house fire, although oxygen and combustible materials are contributory, it is often the discarded cigarette butt that is deemed "The Cause." This process of selection necessitates discerning genuine causes from mere background conditions.[11] To be selected as 'The Cause' is crucial in causal statements like 'x is the cause of y,' as well as in assertions such as 'x caused y,' 'y happened because of x,' or even descriptive forms like 'x burned y.' Each of these statements involves choosing one condition from many potential contributors to represent as the singular cause.

Considering the gap between the fact that causation involves a set of factors jointly leading to a result and the fact that causal statements select one of these factors as "The Cause," it is crucial to first differentiate between two types of questions, which we consider prerequisites for formulating a causal statement:

i. What justifies statements like "c is **a** cause of/condition for e"?

   This is the question that reductionist approaches to causation aim to address. They seek a unified theory of causation that incorporates logical and metaphysical considerations, such as necessity (Mackie, 1965), counterfactuality (Lewis, 1973), or regularity (Davidson, 1967a). This question does not presume singularity; rather, it acknowledges that a cause or condition is typically part of a set of factors that jointly lead to a result.[12]

ii. How is a specific factor selected as "the cause"?

   This process, known as **causal selection**, is guided by different types of considerations, such as moral accountability, or notions of abnormality.[13] Recognizing the complex nature of causal-

[10]It goeas back to (Hume, 1779), who after making the distinction between Ideas and Relations in Part I of *A Treatise of Human Nature*, in Part III section II, he introduces the notion of causation as a *relation* between a cause and an effect.

[11]For a comprehensive discussion on causal selection, refer to Hart (1959), Mackie (1965, 1974), White (1967), Hesslow (1983), Einhorn (1986), Hilton (1986), and Cheng (1991).

[12]Lewis (1973), for example, emphasizes that his analysis of causation in terms of counterfactuality involves *a* cause and not *the* cause. See also Bar-Asher Siegal and Boneh (2020), for linguistic applications of the differences between the indefinite and the definite descriptions of causes.

[13]Diverse criteria for determining causal selection have been proposed, including conversational pragmatics relative to moral responsibility (Driver, 2008), motivational bias (Alicke, 2011), and conversational considerations concerning accountability (Samland, 2016). The primary factor, as argued in recent works, is the salience of abnormality (Knobe, 2010; Halpern, 2015; Blanchard, 2017; Icard, 2017). An event is characterized as "abnormal" if it contravenes a norm, be it statistical (Hilton, 1986) or prescriptive (Sytsma, 2012). Additional situational features posited as relevant for the selection of The Cause include the temporal sequence of potential causes (Einhorn, 1986; Ngbala, 1995; Henne, 2021). Some philosophers have concentrated on conversational principles, premised on assumptions about the state of knowledge and interests of the causal judgment seeker (Beebee, 2004).

ity requires the development of comprehensive models that accurately capture the intricate relationships among the various factors leading to a particular outcome.[14]

We propose adding a third dimension to the considerations that lead to a causal statement. To illustrate, consider an observer who sees a door open following a button press. She might describe the event as "Sam caused the door to open," which, while similar, carries a different inferential weight than simply stating "Sam opened the door." Building upon the insights of Bar-Asher Siegal et al. (2021), we introduce a novel concept: the Causative Construction Selection (CC-selection) process. This process works in conjunction with causal selection and involves choosing a causative construction that precisely characterizes the sequence of events. Each construction has unique licensing parameters that influence how conditions are framed as causes. Unlike causal selection, which focuses on identifying the condition that serves as the cause, CC-selection emphasizes the linguistic choices available to speakers, which subsequently shape their causal judgments. This shift leads us to answer a third question:

> Considering each of the relevant conditions, can it be encoded as the cause in a statement articulated through a specific causative construction, to describe a particular state of affairs?

CC-selection does not aim to identify the specific cause selected by a speaker from a range of conditions; rather, it delineates the linguistic expressions a speaker *may* use when describing the connection between particular events or actions and their result.

By adopting this methodology, we explore the criteria that govern the appropriate use of specific causative linguistic expressions, thereby revealing their truth conditions from a semantic perspective. The CC-selection process identifies, from a range of potential conditions, those that qualify as the cause within a particular causal construction. This process is fundamental to establishing the truth conditions for these expressions. The logic underpinning this methodology can be outlined as follows:

- The constraints represent the necessary conditions that must be met to appropriately employ the relevant causative linguistic expressions.

- The use of these expressions entails or presupposes that these conditions have been met, to describe a particular state of affairs.

- Semantically, this implies that these conditions constitute the truth conditions for the expressions.

As noted in the introduction, the primary contribution of this paper lies in defining the truth conditions of causal statements. By examining the interactions among elements within causal models, we deepen our understanding of how causation is encoded and interpreted in linguistic structures. This analysis underscores the inadequacies of the primitive predicate CAUSE in providing a comprehensive account of linguistic causation. To address this, we propose Structural Equation Modeling as a framework for capturing the multifaceted nature of causality. This framework conceptualizes causality as a network of interrelations among multiple factors, offering a rigorous tool for analyzing processes such as causal selection and CC-selection.

[14]Experimental paradigms for understanding the process of causal selection typically confront participants with scenarios featuring two antecedent events (A and B), succeeded by a target event. Participants are then prompted to evaluate statements such as "Event A caused the target event" and "Event B caused the target event." Nevertheless, research on causal selection often inadvertently neglects critical variables due to variations in the constructions used. Linguistic elements can indeed exert a profound influence on causal judgments, interacting with external causal relations and cognitive perceptions. A noteworthy departure is found in the work of Wolff (2003), who delved into the cognitive correlates of "direct causation" within linguistic contexts, investigating the impact of specific constructions. This is further complemented by the studies of Beller et al. (2020) and De Freitas et al. (2017).

2.2 Structuring a Model for Causation

When reasoning about causality, even in seemingly simple systems, we naturally draw on a set of intuitions about how multiple factors combine to produce a single outcome. As we mentioned, when we consider a scenario where a house fire occurs, we intuitively recognize that a fire may result from various factors—a lit match, the presence of flammable material, and an oxygen-rich environment—that interact in a many-to-one relationship to produce the event. This cognitive process suggests the existence of an underlying mental framework that organizes and represents causal information. We refer to this framework as a "causal model."

This paper aims to formalize such intuitions by establishing a framework for modeling the semantics of causal statements and using it to understand the factors that license the use of specific causal constructions in natural languages. To this end, we apply the Structural Equation Modeling (SEM) framework to represent causal structures. Structural causation models articulate the relationships between propositions representing states of affairs, enabling us to map and analyze the causal dependencies between events and outcomes.

These models extend beyond simple pairings of events linked by predicates such as CAUSE; they provide a structured representation of the relationships among propositions and their truth values, reflecting our shared knowledge of the world. This includes an awareness of the potential interrelations among propositions and the concept of possible worlds that align with natural laws and principles. Speakers draw on this shared causal knowledge to infer the configurations of states of affairs and their interconnections. Thus, identifying a specific set of conditions inherently implies the existence of their predicted effect.

We turn now to illustrate with an example of SEM, as represented in Figure 1. Considering the automated door scenario from the beginning of the paper, we identify certain propositions, denoted A-E, that serve as variables paired with their truth values. The structure of our model dictates that some propositions are contingent on others, a relationship detailed in structural entailments F-G. These entailments emerge not from logical reasoning or lexical meanings, but from familiarity with the door's operational mechanics, essentially encapsulating what is typically known as "world knowledge."

Variables are categorized into two distinct groups: **Exogenous**, or "external," variables, which are independent of other variables, and **Endogenous**, or "internal," variables, which derive their values from the exogenous variables to which they are causally connected. In the context of the door example, variables [A, B, D] are exogenous, while [C, E] are endogenous.

A notable point in our model is the asymmetrical nature of the structural entailments. As outlined in lines (F) and (G) of Figure 1, certain conditions must be met for the endogenous variables to assume specific values (e.g., Button=1; Circuit=1). The door's opening (Door opens=1) is contingent on all these conditions being satisfied (Circuit=1, Electricity=1, Lock=0). Therefore, the model dictates that if any variable within the pertinent set (G) deviates from these values (Circuit=0, Electricity=0, or Lock=1), the door will not open (Door opens=0). This emphasizes that **there are invariable dependencies linking specific values of certain variables to particular outcomes.**

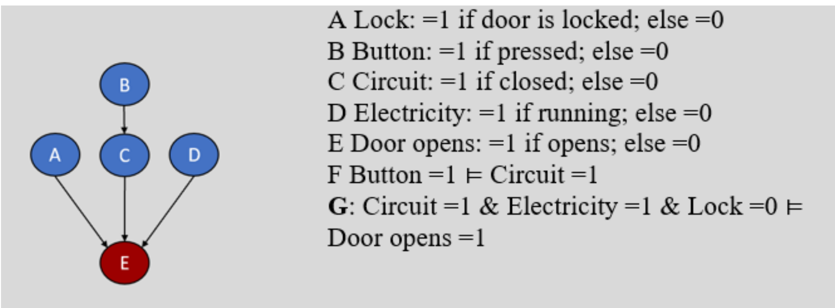

Figure 1: Directed graph and Variable and Structural equations

In this framework, we define a **condition** as any variable that holds relevance—i.e., it causally influences—the value of another variable. "Relevance" here implies that the value of the condition can be invoked, either on its own or alongside other conditions, to deduce the value of a dependant variable. This allows us to pinpoint both the immediate causal precursors to a proposition and the direct nature of their dependencies.

The SEM framework also affords a qualitative depiction of dependencies through directed acyclic graph models, as shown in Figure 1. Here, nodes are associated with variables and arrows signify the directionality of the dependency, with the value of a source node influencing the values of the nodes it directs to. The variables and entailments are mirrored by the graph in Figure 1, making clear which nodes act as "conditions." The graph delineates the variables exerting a direct effect on the endogenous variables, albeit without detailing the exact nature of the dependency.

We argue that these models offer a foundational framework for analyzing the semantics of sentences containing causal elements, providing the structural framework for interpreting causal expressions. As we will demonstrate, the truth of causal statements is contingent on the occurrence of specific events, represented as variables with particular values in a real-world context. At the same time, a causal claim linking cause and effect relies on the existence of a causal model that establishes the dependencies among states of affairs.

3 Formal definitions

3.1 Causal structure and causal model

This section introduces a formal exposition of the causal framework under consideration.[15] Building upon the intuitive visualizations provided by the graphs in Figure 1, Definition 1 defines the interrelations among variables (represented as nodes) within a model.

> **Definition 1 (Causal Structure)**: A causal structure over a set of variables $\mathcal{P}$ is a directed acyclic graph (DAG) in which each node corresponds to a unique element of $\mathcal{P}$, and each link represents a direct functional relationship between the corresponding variables.[16]

A causal structure serves as a foundational blueprint for constructing a "causal model," which precisely defines how each variable is influenced by its parent variables within the DAG. A causal structure (such as the one represented in Figure 1) is the structural basis for a "causal model" in Definition 2. Accordingly, such graphs only identify the variables that have influence on other variables, they do not specify the exact nature of the dependency.

> **Definition 2 (Causal Model)**: A causal model is defined as a pair $M = \langle D, \Theta_D \rangle$, where D signifies a causal structure, and Θ_D represents a set of parameters compatible with D.

[15]We would like to outline the novel contributions of our study in the context of the causal modeling literature. We build upon existing frameworks to define the dependency between Causal Structure and Causal Model (Definitions 1-2), leveraging the foundational work of Pearl (2000) as a guiding inspiration. Our model builds on the approach introduced by Schulz (2007), adapting its framework to incorporate a function that assigns values to variables. These initial steps are not novel per se and we acknowledge the significant groundwork laid by previous scholars. However, contrary to Schulz, our approach does not enforce a strict separation between exogenous and endogenous variables, applying an interpretive function universally across any given set of antecedents. Additionally, we utilize Schulz's framework for presenting the model through truth tables (Table 2), enhancing its clarity and utility in practical analysis. The primary innovation of this paper lies in Definitions 3–7, which integrate the theoretical insights of Mackie (1965) into the SEM framework. This synthesis introduces entirely new formal definitions within causal models, specifically designed to advance the analysis of the semantics of causal expressions. These definitions include formal accounts of necessity and sufficiency. While earlier accounts of necessity and sufficiency with respect to causal models exist (e.g., Baglini and Francez (2016); Nadathur and Lauer (2020)), the definitions proposed here are new and distinct in their formulation and serve a different analytical purpose, as will become evident in Section 3.2.

[16]As noted in the previous footnote, Definition 1 is based on Pearl (2000), the following is his explanation of the functional relationship: "Here we assume that Nature is at liberty to impose arbitrary functional relationships between each effect and its causes and then to perturb these relationships by introducing arbitrary (yet mutually independent) disturbances. These disturbances reflect "hidden" or unmeasurable conditions and exceptions that Nature chooses to govern by some undisclosed probability function."

i. Let ψ denote a node within the causal structure D ($\psi \in \mathcal{P}$).
ii. Let Γ represent a set of nodes on which ψ causally depends within D (i.e. a set of variables for which ψ is a descendant).
iii. $I : \Gamma \rightarrow \{0, 1, u\}$ is an interpretation that maps the set of variables Γ of M to values.
iv. The parameters Θ_D assign a function f that determines the value of a variable ψ based on the interpretation of its antecedents Γ in M within the structure D:
$$[[\psi]]^{M,I} = f(\psi)^{M,I}$$

This type of causal model framework is highly versatile, accommodating a range of variable types, including those representing physical quantities such as temperature readings on a thermometer or a vehicle's velocity. However, for the purposes of this discussion, which focuses on the meaning of linguistic expressions, we narrow our focus in Definition 2 to variables represented as a set of propositional letters $\mathcal{P}$, whose values correspond to truth valuations.

Definition 2 elaborates on how a causal model, given its structure, interprets variables (i) based on the antecedents' values (ii-iii), through an interpretational function f. It specifies how the values of descendant variables are determined. The model—conceptualized as a set of equations or entailments, as illustrated in Figure 1—encodes the information necessary to define the function f that assigns values to endogenous variables, contingent on the interpretation of the antecedents of a given variable. Crucially, in our approach, the interpretation function I operates uniformly across all antecedents, irrespective of whether they are exogenous or endogenous variables.

From a linguistic perspective, a causal model encapsulates knowledge regarding the interplay between different states of affairs. It reflects a speaker's understanding of how real-world events are interconnected and the expected outcomes when certain conditions prevail. In the context of causal statement semantics, we posit that such knowledge underpins the basis for validating speakers' linguistic judgments.

We adopt a three-valued logic system, assigning the values 1, 0, and u (undefined), where "undefined" holds epistemic significance (cf. Kleene, 1952, 335). Importantly, this designation does not indicate ambiguity but rather the absence of definitive knowledge regarding a variable's value. The use of a three-valued epistemic logic aligns with our goal of capturing and modeling speakers' causal knowledge within the SEM framework. Epistemically, causal knowledge relates to the potential causal deductions speakers can make, notwithstanding incomplete information about all prevailing conditions. For instance, if an individual knows the door's opening button has been activated and electricity is coursing through the pertinent circuit but lacks information about the door's lock status, she cannot conclusively determine whether the door will open. Conversely, if she is aware the door is locked, ignorance regarding the button's status still permits her to deduce the door remains closed. Thus, exogenous variables derive their values autonomously, informed either by general world knowledge or specific situational awareness. Endogenous variables, meanwhile, may have their values directly ascertained or inferred within the model, contingent on the causally linked nodes' truth values.

Truth tables serve as an effective means to represent causal linkages, delineating a proposition's value contingent on its descendant propositions' truth values. Table 2 showcases a partially completed model for the automatic door scenario illustrated in Figure 1, elucidating the effect's assigned truth value relative to various causal factors (button activation, circuit continuity, electrical flow, and lock status).

The entailments depicted in Figure 1 outline the necessary conditions for the door to be open. For instance, Table 2 illustrates that under specific conditions—circuit continuity = 1, electrical flow = 1, and lock status = 0—the door opens (Door opens=1). It is crucial to highlight that in scenarios where the state of any causal factor remains undefined (represented as "u"), the outcome of the effect, namely the door's opening, cannot be determined with certainty (Door opens=u). Comprehensive knowledge of all three requisite conditions is essential to assert that the door will open. This is predicated on the assumption within this framework that the door is closed by default; thus, failure to meet these conditions implies the door remains closed. Therefore, evidence of a disrupted circuit, no electricity, or a locked door is sufficient to deduce the door's closure—even without complete information on all three conditions (for example circuit continuity = 0, electrical flow = 1, and lock status = u—entails that the door is close (Door opens=0)). Understanding the asymmetry between the evidence needed to conclude the door is open versus what is necessary to ascertain its closure is fundamental to appreciating the concept of sufficiency in causal analysis, that will be discussed below (cf. Glass 2023).

Furthermore, the absence of information regarding the value of an exogenous variable does not preclude the inference of its antecedents' states when independent knowledge about an intermediary node is available. For instance, knowing that the circuit is active eliminates the need for direct information about the status of the button.

Button	1	1	1	1	0	0	0	u	u	u	u	u	u	u	u	u	1
Circuit	1	1	1	1	0	0	0	1	1	1	0	0	0	u	u	u	1
Electricity	1	1	0	0	1	1	0	1	1	0	0	0	1	1	0	1	u
Lock	1	0	1	0	1	0	0	1	0	0	0	1	0	1	1	0	0
Door-open	0	1	0	0	0	0	0	0	1	0	0	0	0	0	0	u	u

Table 2: Automatic door example

Up until now, terms such as "sufficient set" and "necessary conditions" have been used informally and intuitively. The next subsection provides formal definitions of these notions within the SEM framework.

3.2 Defining sufficiency

3.2.1 Mackie's INUS conditions

In order to structure a causal model for the interpretation of causal statements, it is necessary to examine how structural models capture dependencies among states of affairs. These models represent relationships between propositions, corresponding to variables and their truth values, and how sets of variables can influence the truth value of another variable. This framework emphasizes that a given effect often depends on a complex interplay of conditions. A set of formal definitions is essential for precisely articulating these relationships:

i. A definition for a proposition (a variable assigned a value) that functions as a condition within the model, influencing the values of another variable.
ii. A definition for the manner in which such a condition interacts with other conditions, either collectively contributing to the same outcome or independently capable of producing that outcome.

Various philosophical accounts of causation have proposed distinct frameworks to address these issues and provide such definitions.[17] This paper adopts Mackie's (1965) *characterization* of causation as the foundation for our analysis. Within this framework, a condition in a causal model is considered **necessary** if it contributes to the occurrence of a specific outcome (i). Such a condition interacts with other conditions to form a **sufficient** set for that outcome (ii). To illustrate, we employ the door-opening scenario depicted in Figure 1. This example demonstrates a case in which the desired outcome—opening the door—depends on the joint presence of multiple conditions. While each factor is individually **necessary**, it is the collective presence of these conditions that constitutes a **sufficient** set for the result.

Our framework elaborates on Mackie's concept of INUS conditions. In his original formulation, each condition is individually Insufficient but Necessary, and their combination is Sufficient to bring

[17]See Mill (1884), particularly Chapter 8, for foundational contributions to this line of inquiry. In linguistics, Hobbs (2005) explored related ideas through the notion of "causal complexes," in a different framework. However, it is beyond the scope of this paper to explain why he cannot account for the data discussed in the current paper.

about the outcome. Additionally, Mackie observed that these sets are collectively Unnecessary, a perspective that accommodates cases where multiple causal pathways lead to the same result. For instance, alternative methods for opening a door reflect distinct pathways within the causal model. Figure 2 extends this discussion by introducing a new variable (A) and an inference relation (I) to illustrate the importance of accounting for multiple causal pathways within the model.

With the addition of INUS sets, our proposed approach extends beyond accounting only for *actual causation*, which seeks to explain causal attribution in specific events. Instead, it aims to represent **broader causal knowledge**, capturing how various factors **might** interact to bring about outcomes. This broader causal knowledge is encoded within causal models, enabling the representation of alternative sets of sufficient conditions and the *possible* interactions among conditions that generate effects.

The extension to INUS sets also enhances the model's capacity to reflect real-world causality by incorporating potential causal pathways and detailing the multiple configurations through which effects can arise. Following VanderWeele and Robins (2009), Figure 2 illustrates with a circle conjoined conditions within the causal graphs, demonstrating that the outcome depends on the interaction of several factors.[18]

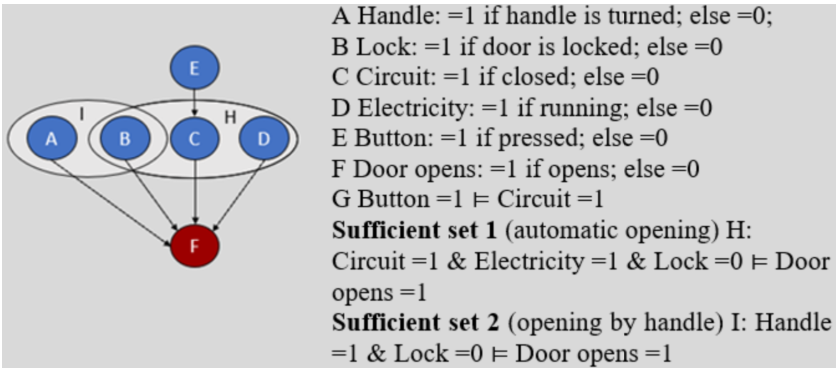

Figure 2: Graphical models of two sufficient sets of conditions for an effect F

The application of Mackie's INUS conditions provides two key contributions to our approach. First, it allows us to define a condition within the model, understood from the perspective of Structural Equation Modeling (SEM) as a relevant variable or node within the graph. Second, it offers a way to conceptualize how various conditions jointly contribute to producing an outcome, as illustrated in Figure 2. In causal models, these conditions are represented as variables, enabling us to formalize the relationships among them. This "translation" into variables allows us to characterize the dependencies between conditions as inferential relations between propositions, where the values of these variables correspond to truth valuations.

Furthermore, Mackie's perspective significantly contributes to the discourse on causation by asserting that sufficiency relates to a collection of causal factors rather than individual conditions. This view diverges from the prevailing notion of causal sufficiency in the literature of philosophy, cognitive sciences, and linguistics, which typically frames sufficiency as a direct relationship between a single condition and its resultant effect.[19] Commonly, necessity and sufficiency are informally defined as follows:

Necessity: If C were not to occur, E would also not occur.
Sufficiency: If C were to occur, E would also occur.

[18]As in the case of Mackie, being part of a set that is not necessary for the effect is not a requirement for the definition of "a cause." There can be cases where a singular sufficient set alone can precipitate the result. Also, Figure 2 does not label all sufficient sets, as EBCD is also one. This observation will become relevant later in our discussion.

[19]Notably, efforts have been made to formalize the distinction between necessity and sufficiency, capturing sufficiency in terms of production Pearl (2000, 286) and exploring various formal definitions relevant to linguistics (Baglini and Francez (2016); Martin (2018); Bar-Asher Siegal and Boneh (2019)) and to cognitive sciences (see Mandel (2003), and more recently Icard et al. (2017)).

These definitions require the context of background conditions for their applicability (cf. Nadathur and Lauer (2020)), highlighting a focus on *actual* causation rather than merely describing *conditions within a causal model*.

Our approach, drawing on Mackie's framework, conceptualizes sufficiency not as a property of isolated conditions but as a relationship among sets of propositions. This perspective is essential for identifying and distinguishing between various sufficient sets. It shifts the focus from the mere sequential occurrence of events to the broader structural organization of the model, offering a more nuanced understanding of the causal architecture underlying the world.

Building on the formal definitions of causal structure and causal model (Definitions 1–2), we will now formally define necessary conditions and sufficient sets. Establishing these definitions will enable us to address questions of causal selection and CC-selection, both of which involve the selection of specific conditions from the broader set of factors contributing to an outcome. This formal framework will allow us to account for the semantics of causal expressions, as CC-selection is integral to their interpretation in our approach.

3.2.2 Formal definitions to necessity and sufficiency

The goal is to define what constitutes an INUS condition. To accomplish this, we must address the inherent circularity present in its traditional definitions, where each necessary condition is identified as part of a sufficient set, itself defined by the necessary conditions required to produce the effect. To overcome this circular reasoning and establish a robust framework for analyzing causal relationships, an external anchor is indispensable. This anchor provides a reliable foundation for defining INUS conditions. To achieve this, we first delineate a set of propositions that describe a situation, as outlined in Definition 3. Following this, we introduce the concept of "causal relevance" in Definition 4. Causal relevance requires a degree of causal dependency, though it does not necessarily entail sufficiency.

> **Definition 3 (Situation)**: A set of propositions, pairs of variables Σ in $\mathcal{P}$ and their 0/1 valuation, is a situation S.

A situation can be conceptualized as a set of possible worlds wherein a specific set of conditions obtains.[20] As previously noted, causal relationships connect distinct states of affairs, depicted through propositions. This connection is demonstrated in truth tables (akin to Table 2), where a column's set of truth values corresponds to the variables. Consequently, the notion of situations facilitates the illustration of connections between conditions (denoted by variables and their truth values) and their resultant effect (also denoted by a variable and its respective value). This foundation allows for the subsequent definition of causal relevance in Definition 4.

> **Definition 4 (Causal relevance)**: A situation S (consisting of a set of pairs of variables Σ in $\mathcal{P}$ and their values) of M is causally relevant for a certain value of the variable ψ, if the following conditions hold:
>
> i. Let I be the interpretation of the set of all variables $\mathcal{P}$ except ψ in M in situation S, with the following constrains:

[20] In Stalnaker's framework, a proposition is conceptualized as a set of possible worlds in which it is true (Stalnaker 1984; Stalnaker 1987). This set represents the worlds where the truth conditions of the proposition are satisfied, delineating truth conditions across possible worlds and classifying them based on the truth values of propositional variables. Similarly, a situation—a set of variables and their values, which correspond to propositions—can be understood as the set of possible worlds where a specific set of conditions holds. Truth tables, such as those shown in Table 2, can thus be seen as representing variations across relevant possible worlds, with each column corresponding to a distinct set of worlds.

ii. $I : \Sigma \rightarrow \{0, 1\}$.
iii. $I : \Sigma^c \rightarrow u$, where Σ^c represents the complement set of Σ in $\mathcal{P}$.
iv. $[[\psi]]^{M,I} = [[f(\psi)]]^{M,I} \in \{0, 1\}$.

Causal relevancy denotes the connection between a situation, represented by a set of variable-value pairs, and another situation represented by a single variable and its corresponding value. It signifies that an inference can be made from the collective interpretations of the variables to ascertain the value of the other variable. Specifically, condition (iv) states that the value of ψ is well-defined (1 or 0), when I assigns certain values of 0 or 1, to the set of variables Σ, which constitutes the situation S (condition ii), and assigns value u to all members of the complementary set of Σ in $\mathcal{P}$ (condition iii). The combination of these conditions ensures that all relevant factors necessary to determine the value of ψ are included within the situation S.

However, it may be the case that not all pairs within the situation S possess equal significance. For example, in our illustrative example, the set $\{ < Handle, 1 >, < Lock, 0 >, < Electricity, 1 > \}$ is causally relevant for $< Door\ open, 1 >$; that is, if situation S prevails, the door must be open. However, our analysis seeks to recognize that the pair $< Electricity, 1 >$ is superfluous in this context. Specifically, this element is not required for the outcome in question, as a similar effect ($< Door\ open, 1 >$ - the door being open) is observed in situation S', which only differs in the value assigned to the variable Electricity, including the pair $< Electricity, 0 >$ instead. Consequently, the focus shifts towards delineating the concept of causal necessity in Definition 5, aiming to establish that within the situation S, the pairs $< Handle, 1 >$, $< Lock, 0 >$ are necessary, whereas the pair $< Electricity, 1 >$ is not. In the spirit of Mackie, our aim is to be able to identify the INUS conditions, those that are individually necessary for the occurrence of a specific effect within a given situation (a set of conditions).

Building on the concept of **causal relevancy** and the definition of **a situation**, we can now proceed to formally define (Definition 5) causal necessity as a relationship between two valued propositions, namely a cause and its corresponding effect, in the context of a given situation.

Definition 5 (Causal necessity): In a given situation S, the variable χ and its corresponding value ($S\chi$) are causally necessary for a specific value of ψ (0 or 1) if the following conditions are satisfied:

i. The situation S is causally relevant for the value of the propositional variable ψ.
ii. $I : \Sigma \rightarrow 0, 1$ is an interpretation of the set of variables Σ of M in situation S, and $I' : \Sigma \rightarrow 0, 1$ is an interpretation of the set of variables Σ of M in situation S' (The relationship between situation S and situation S', is characterized by the following two constraints (iii-iv)).
iii. $S\chi \neq S'\chi$, indicating that the value of χ in situation S is different from its value in situation S'.
iv. The cardinality of the set J, the symmetric set difference between S and S', denoted as $J = S \triangle S'$, is 2. This means that the set J contains two elements, representing the pair χ and its corresponding value, which differ between the two situations.
v. $[[\psi]]^{M,I} \neq [[\psi]]^{M,I'}$, indicating that the interpretations of ψ in situations S and S' respectively, yield different truth values.
vi. There is no interpretation I'' of the set of variables Σ where $S\chi \neq S''\chi$ and $[[\psi]]^{M,I} = [[\psi]]^{M,I''} \neq u$.

Condition (i) allows us to focus in the model on the propositions that are part of the situation S, disregarding the values of other variables. This is based on the concept of causal relevancy as defined in Definition 4, which ensures that the values of all other variables in the model are irrelevant if the casually relevant set in focus is realized, therefore they are assigned the value "u". Additionally, due to condition (i), the values of $S\chi$ and $S'\chi$ in condition (iii) are restricted to 0 or 1, excluding the possibility of "u".

Definition 5, with conditions (i-vi), enables us to precisely examine the influence of one variable χ and its corresponding truth value on the value of another variable ψ, given the situation S. The purpose of conditions (ii-v) is to compare two situations that differ minimally (conditions ii-iv), with only a single variable's truth value varying between them. Noting the disparity in the interpretation of ψ between these two situations in model (v), we can conclude that the value of χ in the causally relevant situation is indispensable for the specific interpretation of ψ. Hence, it becomes apparent that the pair consisting of χ and its truth value is non-redundant in the particular causally relevant context (S) for ψ.

Revisiting our ongoing example, the set $\{ < Handle, 1 >, < Lock, 0 >, < Electricity, 1 > \}$ is identified as causally relevant to $< Door\ open, 1 >$. Within this set, the pairs $< Handle, 1 >$ and $< Lock, 0 >$ are classified as necessary conditions for the door to be open, as altering the values of these variables results in the door remaining closed. In contrast, $< Electricity, 1 >$ is not considered necessary, since changing its value to $< Electricity, 0 >$ does not affect the state of the door being open.

Considering the constraint in (vi), this analysis draws on von Wright's (1974) observation regarding the interdefinable relationship between sufficient and necessary conditions: *Necessary*(p, q) $\equiv$ *Sufficient*($\neg p$, $\neg q$). Within our framework, this principle captures the idea that if a condition is deemed necessary within a set of conditions for a specific outcome, then its absence is sufficient to prevent that outcome from occurring. In the context of the electronic door example, the manual scenario illustrates this relationship: Handle=1 is a necessary condition for Door open=1, but it is also required that the door is unlocked (Lock=0). Conversely, if Handle=0, regardless of the value assigned to Lock, the outcome will invariably be Door open=0. This encapsulates the meaning of condition (vi). When restricted to the set of variables Σ, if the value of χ differs from that in situation S ($S\chi$, which is a necessary condition for a specific value of ψ), then the value of ψ must necessarily invert.

Definition 5, which pertains to *necessity*, shares similarities with David Lewis's (1973) definition of causation framed through *counterfactuality*. This connection ties back to our earlier discussion that situations can be viewed as formal representations of possible worlds. Consequently, we can consider the differences between situations S and S' in Definition 5 as analogous to closely comparable possible worlds. Specifically, a change in the value of one variable leading to a different truth value for the effect variable precisely mirrors Lewis's approach to counterfactuality, which involves comparisons between similar possible worlds.

While a detailed comparison between this approach and Lewis's framework of causation is beyond the scope of this paper, it is crucial to outline the key distinctions. Lewis's model primarily addresses what we called earlier *actual causation*, focusing on assessing similarities between possible worlds and the actual world (or any other specific possible world). In contrast, our definition of necessity explores the relationships between conditions within the models, irrespective of the occurrence of any event. Thus, in Definition 5, necessity is conceptualized by comparing possible situations, deliberately omitting direct references to the actual world.[21] This approach not only emphasizes the importance of evaluating causal factors in relation to other causal elements but also underscores that *necessity* is defined relative to a given situation.

Once we have a definition of causal necessity, and of situations it is possible to define (Definition 6) a sufficient set, all of whose members are necessary.

Definition 6 (Sufficient set): A situation S is defined as a sufficient set for a certain interpretation of ψ, if the set of variables Σ whose values are defined in S (only 0-1 valuations) is causally relevant for this interpretation of ψ, and every member of S is causally necessary for this interpretation of ψ.

[21]See also Halpern and Pearl (2005) for a conceptual comparison between SEM and Lewis's definition of actual causation in counterfactual terms.

The following section explores the application of the concept of a sufficient set within our framework, specifically illustrating how sufficient sets are identified in our case study involving the opening of a door.

3.2.3 *A* sufficient set vs. *a completed* Sufficient set

According to Definition 6, a situation qualifies as a sufficient set if it meets two criteria: 1) it is causally relevant (per Definition 4) to a certain interpretation of the effect-variable, and 2) every component within it is deemed causally necessary (as specified in Definition 5) for that interpretation. Causal relevance ensures that, once all the specified conditions in the relevant situation are met, the effect must invariably occur. Given that each condition is necessary, they all contribute to the outcome.

Importantly, Definition 6 does not stipulate that a sufficient set must include every necessary condition for the effect's occurrence. Therefore, the essence of a sufficient set is not the list of conditions that, when combined, **produce** the effect. Instead, it is a set of conditions such that, once confirmed to have occurred, one can confidently **deduce** that the effect has also materialized. This subtlety in the definition allows for the existence of multiple sufficient sets for the same result within the same model, each varying slightly from the others. What they all share is that the occurrence of these conditions guarantees the same result. In the context of the electric door example shown in Figure 1, three distinct scenarios are identified as sufficient sets for the variable F taking the value 1:

- Sufficient Set 1 $\{< E, 1 >, < C, 1 >, < D, 1 >, < B, 0 >\}$
- Sufficient Set 2 $\{< C, 1 >, < D, 1 >, < B, 0 >\}$
- Sufficient Set 3 $\{< E, 1 >, < D, 1 >, < B, 0 >\}$

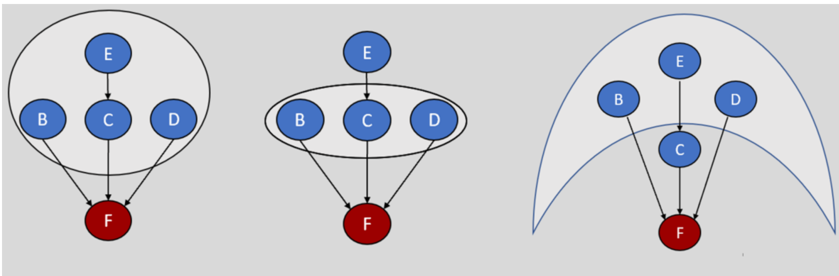

Figure 3: The various sufficient sets

This represents a significant deviation from Mackie's viewpoint, where he considers an INUS condition as part of a sufficient set that encompasses all the necessary conditions to produce (together) the outcome (referred to as **"a completed sufficient set"**).

Intriguingly, there are situations that consistently involve various sufficient sets. The following discussion aims to delineate these cases with greater specificity:

According to Definition 5, causal necessity establishes a binary relationship between two conditions (valued variables) in two distinct situations, where the second situation encapsulates a single condition. This relationship can be represented as CauseNec$(< \chi, S\chi >, < \psi, R\psi >)$, indicating that χ within situation S is causally necessary for a specific valuation of ψ in situation R.

Notably, this relationship is characterized by its transitivity. If a condition ϕ in situation N is necessary for the pertinent valuation of χ in situation S (CauseNec$(< \phi, N\phi >, < \chi, S\chi >)$), and if CauseNec$(< \chi, S\chi >, < \psi, R\psi >)$ is valid, then CauseNec$(< \phi, N\phi >, < \psi, R\psi >)$ is likewise affirmed.

This transitivity follows directly from the nature of causal necessity. When S is a subset of N, if CauseNec$(< \phi, N\phi >, < \chi, S\chi >)$ holds, it indicates that altering the valuation of ϕ in situation N would lead to a different valuation of χ in situation S. Similarly, if CauseNec$(< \chi, S\chi >, < \psi, R\psi >)$ is valid, it implies that changing the valuation of χ in situation S would result in a different valuation of ψ in situation R. From these premises, it logically follows that CauseNec($< \phi, N\phi >, < \psi, R\psi >$) must also hold.

The inherent transitivity within the concept of causal necessity unveils the insight that a particular outcome can emerge from various sufficient sets. More explicitly, it reveals that one sufficient set may exist as a subset within another for the identical effect. In a scenario where the aforementioned necessary relations are established, the following deduction is applicable:

i. If the situation S constitutes a sufficient set for a specific value of the proposition ψ; and
ii. Situation N encompasses situation S as a superset ($N \supset S$),
iii. Provided that the set difference between N and S (elements present in N but not in S) consists of a single element, with ϕ being that element, and given that $N\phi$ is a necessary condition for $S\chi$ (CauseNec($< \phi, N\phi >, < \chi, S\chi >$)),
iv. It follows that situation N also serves as a sufficient set for the same value of proposition ψ.

The phenomenon of having different sufficient sets, where one set is a subset of another, directly follows from the nature of Definition 6. The realization that a superset of a sufficient set can also serve as a sufficient set for the same outcome has significant implications, particularly in the context of deterministic chains. In such cases, a sequence of deterministically linked conditions contributes to the outcome, often alongside other sets of conditions. The most comprehensive sufficient set includes all the conditions in the deterministic chain.

However, according to Definition 6, various subsets of these conditions may also qualify as sufficient sets for the same outcome. Each subset includes all necessary conditions and at least one condition from the deterministic sequence. The inclusion of this particular condition ensures the fulfillment of all subsequent linked conditions, thereby satisfying the relationships specified in (i–iv). These subets are causally relevant to the outcome, as the fulfillment of all their conditions guarantees the result, and each element within the subset is necessary. Consequently, they meet the criteria for a sufficient set as outlined in Definition 6.

Considering the electric door scenario in Figure 1, activating the button is a necessary condition for closing the circuit (both are links in a deterministic sequence). Given that the circuit's activation is a necessary condition for opening the door, pressing the button is also a necessary condition. This analysis aligns with the initial observation of this section, identifying three distinct scenarios that form sufficient sets for achieving the value 1 of F, as demonstrated in Figure 3.

The relevance of this observation will become evident in the linguistic analysis of inferential relationships between sentences, as discussed in Section (3). In addition, Section 5.6 revisits the notion of the **completed sufficient set**, providing a precise definition and demonstrating its role in capturing the semantics of causative constructions.

Having established the foundational elements of a causal model within the framework of SEM and defined the notion of a condition in the context of sufficient sets, we are now positioned to address the primary objective of this paper: developing a framework for analyzing the semantics of causal statements that involve CC-selection.

4 Application in natural languages: The direct causation puzzle

The introductory section of this article outlined its primary objective: to explore the nuances of meaning in causative constructions and to develop formal models that accurately capture these subtleties. Building on the foundation established in earlier sections, which examined the formal structure of causal models, the paper now turns to substantiate the proposed framework. This study employs a unified approach to causation, suggesting that various linguistic expressions of causality can be effectively represented using SEMs, enhanced by the integration of construction-specific entailments and pragmatic inferences. Furthermore, these elements act as constraints that guide the selection of causative factors from among multiple potential conditions—a process referred to as CC-selection, detailed in Section 2.1. A key innovation of this approach is that CC-selection leverages the formal properties of conditions within the model, as outlined in the previous sections. The ensuing discussion focuses on a particularly contentious case in the research on causal language. The analysis of a specific contrast between two causative constructions serves as a proof of concept, illustrating the methodology's effectiveness.

The discussion centers around the observation mentioned in (3), which highlights that change-of-state [=CoS] causative[22] entails the truth of the explicit *cause* sentence, whereas the reverse entailment does not hold.[23] The entailment pattern presented in (3) indicates that the verb *cause* can be applied in a broader range of situations compared to its corresponding CoS causative. As demonstrated in (4), certain situations, including those where Sam's action indirectly or through extended causal chains leads to the opening of the door, are not appropriately described using the verb *open*.

(4) # Sam opened the door.
Context A: Sam asked someone else to open the door.[24]
Context B: Sam opened a window and the resulting gust of wind opened the door.

This difference is often attributed to an additional semantic constraint of direct causation associated with CoS verbs. Specifically, the causative relation requires a spatiotemporally contiguous cause and effect, where the causer directly—frequently physically—manipulates the causee (Shibatani 1976, 31; Pinker 1989, 48) with no intervening third event (Fodor, 1970; Katz, 1970; Rapoport, 1999; Wolff, 2003, inter alia).[25]

While the intuitive notion that a CoS verb describes a causal relation that is 'more direct' than its paraphrase with the word *cause* remains popular, the contiguity-based hypothesis faces both theoretical challenges and well-documented empirical problems (see Neeleman and Van de Koot (2012), for a recent review of the problems with the direct causation hypothesis).

At the theoretical level, while numerous studies rely on an intuitive understanding of direct

[22]For syntactic and semantic justifications supporting the consideration of CoS as a distinct verb category, see Rapaport Hovav and Levin (2002) and Beavers (2013), among others.

[23]According to our description of the inferential relationship between the constructions, the verb *cause* encompasses both direct and indirect causation. This description contradicts the claim made by McCawley (1978) that using the verb *cause* comes with an implicature of indirect cause. McCawley suggests that this implicature arises from a competition between the two constructions (see also Rett (2020) and Benz (2006)). However, empirical studies contradict this claim by demonstrating that the verb *cause* also tends to prefer direct causation, often selecting and favoring more recent conditions over earlier ones (see Henne et al. (2021) and Bar-Asher Siegal et al. (2021)). McCawley's assertion appears anticipated, given that CoS causatives exhibit more limitations compared to the application of the verb *cause.* Consequently, it is presumed that the use of *cause* might invoke implicatures to denote situations beyond the scope of a lexical causative. See Bar-Asher Siegal et al. (2021) for a range of explanations addressing why such implicatures do not emerge.

Furthermore, Martin (2018) proposes that the association of directness with CoS verbs is an implicature, as it can be cancelled. Nonetheless, she argues that even when directness is cancelled, the causal impact is "established from a retrospective perspective only" (p. 118). As we will further explain later (Section 5.5), this aligns with the semantic analysis we propose for CoS verbs.

[24]In this context, the acceptability of the sentence is enhanced when Sam is regarded as the ultimate authority concerning the door's status—whether it is opened or not—and he has authorized its opening. This improvement in acceptability appears to be linked to the concept of foreseeability, as discussed in Section 5.5. In scenarios where a command is issued by an authoritative figure, it becomes foreseeable that the command will be executed accordingly.

[25]Wolff (2003, 3–4) provides a review of the literature on direct causation. In this paper, Wolff, working within the production framework, introduces the no-intervening cause hypothesis, which posits that direct causation requires the absence of intervening causers between the initial cause and the final effect. His experiments suggest that single-clause sentences (lexical causatives) are preferred when (1) the cause and effect are in direct contact, and (2) intervening entities are interpreted as enabling conditions rather than independent causes. However, Wolff's reliance on the notion of "enabling condition," grounded in the productive account of causation (see also Copley et al. (2015)), raises concerns. From our perspective, rooted in the dependency approach, this concept lacks a precise definition, faces significant empirical challenges, and complicates the resolution of counterexamples. For an introduction to the distinction between dependency and productive accounts of causation, see Bar-Asher Siegal and Boneh (2020).

causation, they often lack explicit models to formally represent and substantiate their claims. In fact, capturing the specific felicity conditions associated with CoS verbs presents a serious challenge. It requires the development of a method for modelling potentially intricate causal chains, which is far from trivial (cf. Dowty (1979); Ballweg (1997); and Thomason (2014)). Furthermore, the concept of "directness" lacks a precise metaphysical definition. When examining a chain of events leading to a specific outcome, there will invariably exist potential subevents between a cause and its result, rendering it challenging to determine what truly constitutes a more direct causal link (see, inter alia, Davidson (1967b), for the challenges that come with the individuation of events, especially in the context of directness).

On the empirical level, various observations challenge the hypothesis that excludes the possibility of intermediate causes. For instance, examples shown in (5) (a-b) reveal that CoS causatives may indeed involve intervening causes. It's also important to highlight that any intervening cause mentioned in (5)(a-b) has the potential to be selected as the cause in the relevant scenarios, using the CoS verbs.[26]

(5) a. Opening bus lanes to motorcycles will redden the streets of London with cyclists' blood.
[*opening bus lanes > accidents increase > some cyclists die*]
b. A large fleet of fast-charging cars will melt the grid.
[*many electric cars on roads > many cars charging simultaneously > high electricity demand > heating of electric cables > melting of the grid*]
(from Neeleman and Van de Koot 2012)

However, while certain intervening events can be represented with CoS verbs in (5), the presence of intervening *agents* presents a different scenario altogether. The observation that the acceptability of CoS causatives is disrupted by intervening agent-controlled events was initially made by Katz (1970) and can be illustrated through the following scenario:

(6) # The gunsmith killed the sheriff.
Context: A sheriff has his six-shooter gun faultily repaired by the local gunsmith. As a result, his weapon jams at a critical moment during a gunfight with a bandit and the sheriff is killed.

Katz concludes that "clearly, the gunsmith caused the death of the sheriff, but equally clearly, the gunsmith did not kill him". The point is that while there is a causal chain from the gunsmith's faulty repair to the sheriff dying, the faulty repair nevertheless fails to meet some condition of causal "directness" or "immediacy" required by the CoS causative. What appears to differentiate felicitous examples in (5) from the infelicitous example (6) is the presence of an intervening agent in the latter: that is, the bandit who fires the shot which actually kills the sheriff (see also Cruse 1972 and Shibatani 1976). This would suggest that intervening events do not preclude the licensing of CoS verbs as long as no intervening event is controlled by an agent. But the role of agents in the licensing of CoS verbs is still more complex, as illustrated by contrasting examples like those in (7).

(7) a. The eclipse ended the performance of Beethoven's 5th.
[*lunar eclipse > distracts musicians > performance ends*]
b. While the musicians were performing Beethoven's 5th Symphony, an eclipse coincided, leading to their distraction. After careful deliberation, the conductor made the decision to halt the performance.
?The eclipse ended the concert. The eclipse caused the concert to stop.
[*lunar eclipse > conductor stops conducting > performance ends*]

---

[26] In other words that in this scenario all the following sentences will be acceptable:

> Opening bus lanes to motorcycles/ increase in accidents/ the deth of cyclists will redden the streets of London with cyclists' blood.

Similarly with (5b) and any other chain of this sort.

Notice that both sentences in (7) involve agent-involved intervening causes, but only the latter involves an event that is fully controlled by a volitional agent (the conductor). This suggests that the initial cause of a extended causal chain may be selected as the subject of a CoS verb if and only if no event in that chain is controlled by a *volitional* agent.

Thus, we need to elucidate the asymmetry between CoS verbs and the explicit causative verb *cause* by analyzing their entailment patterns (3) and the empirical complexities associated with causal directness (5)-(7). Consistent with our central thesis that the semantics of causative constructions can largely be captured through constraints on CC-selection, this analysis examines the specific restrictions that CoS causatives impose on the selection of a linguistic subject. Whereas previous research on direct causation predominantly concentrated on identifying a cause within a sequential array of factors (a causal chain), (cf. Dowty (1979, 106)), our approach adopts a wider lens. It includes considering the selection of a cause from a group of contributory factors, or in Mackie's terminology: causally necessary conditions within a set of sufficient conditions. This expanded perspective allows us to integrate the semantic intricacies of CoS verbs and causal chains within the broader framework of causation, accounting for the complex structure of various conditions.

To provide a more detailed illustration of causal selection, let us consider why sentence (8-a) is acceptable while (8-b) is not in the typical scenario of an automatic door opening (when electricity flows uninterrupted through the network):

(8) a. {Sam/pushing the button/the button} opened the door.
    b. #Electricity opened the door.

While the contrast in (8) could still be explained with the notion of direct causation, relying on temporal contiguity (assuming that only the last condition to be fulfilled can be selected as the subject of CoS verbs), let us consider a case in which Sam is on a train which allows the door open button to be held down before the train stops and the door opens as soon as the train completely stops. If Sam presses the button of the door before the train reaches the station, (9-a) is still acceptable. As a matter of fact (9-b) is not straightforward, despite the fact that the door does not immediately open, and it depends on the arrival of the train to the station.[27]

(9) a. Sam opened the door.
    b. ?The train's arrival at the station opened the door.
    Context: Automatic train door with safety delay until the train stops.

We posit that for a participant to be represented as the subject of a CoS verb, it must be involved in an event that causally **ensures** the occurrence of the corresponding effect (denoted by the VP). In essence, the subject must participate in an event where the occurrence of the effect is guaranteed to follow.

The formal contribution of this paper lies in articulating this requirement through the concept of *sufficient sets*. Specifically, it is argued that **the subject must be a component of the condition that completes *a sufficient set* for the effect to occur.**

Moreover, as noted earlier (Section 3.1), it is the epistemic stance of the evaluator (participants in the conversation) that plays a crucial role, focusing on under which conditions the effect's occurrence is deemed guaranteed. This hypothesis leads to three corollaries:

1. Participants in events cannot be selected as the cause in a CoS verb construction if the event they are part of is succeeded by an independent volitional action from an agent, as the effect's occurrence remains uncertain due to the potential for the volitional agent to decide against acting. Thus, the inadmissibility of scenarios such as Katz's sheriff (6), is explained, and the conductor examples (7-b) as well.

2. In causal models characterized by deterministic chains, where each event ensures the subsequent event's occurrence, participants in all events in the deterministic chain can be depicted

[27] This sentence can be improved somewhat under focus for some speakers: *Actually, it was the arrival of the train to the station that opened the door.*

as the subject of a CoS verb construction. The occurrence of these events guarantees the result, which clarifies the acceptability of examples provided by Neeleman and Van de Koot (5)), alongside the example involving a distracting eclipse (7-a).

3. Given that the epistemic stance is critical, the occurrence of certain events can be assumed as given. Hence, the selected cause is the last condition whose occurrence was unforeseen. This facet of the hypothesis elucidates scenarios akin to the one where the train's arrival is expected, and the only uncertainty is whether someone will press the button to open the door (9-a). In such cases, the act of pressing the button can be represented as the subject of the CoS construction.

The subsequent section offers a formal presentation of the ideas introduced informally earlier, demonstrating the significance of the formal definitions formulated in Section 3 for representing the truth conditions of causal expressions.

5 Application to causal constructions in natural languages

5.1 Causal Models vs. causal statements

The aim of this study is to develop a semantic framework for causative statements expressed through causative constructions and to elucidate how their truth conditions are anchored in causal models. This requires a precise articulation of the relationship between the semantics of individual causal statements and the causal models that underlie them.

Causal models are conceptualized as representations of causal knowledge, capturing how individuals reason about and interpret causal relationships independently of language. SEMs delineate causal connections between variables, facilitating the abstraction of general causal relationships. In contrast, causal statements refer to specific events via inflected predicates, presenting a binary causal relationship between events. Despite these differences, our assumption that the truth conditions of causal statements are based on the causal models conceived by speakers. Examining the link between causal models and statements deepens our understanding of causal semantics and the underlying principles that govern their truth conditions.

The interplay between causal models and causal statements is akin to the distinction between type-causal and token-causal statements. Type-causal statements outline broad causation principles and laws, while token-causal statements delve into specific causal event connections.[28] Similarly, causal models serve as generalizations about the relationships between events, translating into nomological (law-like) connections among variables that symbolize types of events, thereby representing regularities intrinsic to token causation (Hausman 2005, Hausman 1998, Woodward 2003). These models elucidate the dependencies between variable values, facilitating a "translation" into type-level causal assertions. Adopting a scientific perspective, this approach elucidates the relationship between causal statements and causal models. The causal model can be understood as the repository of causal knowledge derived from empirical observations (Pearl, 2000), reflecting the transition from specific (token) causal instances to generalized (type) causal assertions. Truth tables, in turn, provide a systematic represention of generalizations that capture the relationships between types of states of affairs.

Conversely, and crucially for our analysis, which centers on token causal claims, establishing a causal model enables a reversed perspective. In this perspective, the validity of token-causation statements is grounded in nomological (law-based) relations. The truth conditions for specific instances of causation, or token-causation, are defined through type-causal relationships (Tooley 1987,

[28]Type-causal statements encompass widespread instances of causation and general causal laws, as seen in declarations like "The door can be opened by pressing the button when there is electricity and the door is unlocked." In contrast, token-causal statements such as "Sam opened the door" or "Sam caused the door to open" highlight particular causal relationships between events, emphasizing the immediate cause and effect rather than abstract properties or categories.

Hoover and Hoover 2001). Consequently, the validity of a causal statement is confirmed by its congruence with the underlying causal model. For instance, the statement "Sam opened the door" is considered true if Sam pressed a button leading to the door opening, but only if it aligns with a causal model such as the one depicted in Figure 2, where the opening of the door (variable F) is contingent on pressing the button (variable E).

Adopting this approach allows us to interpret the meaning of causal statements through concepts integral to causal models, such as necessity and sufficiency. These concepts, as discussed in Section 3, are firmly rooted within the causal model. Necessity, for example, (as defined in Definition 5) is determined by examining hypothetical situations without direct reference to the actual world.

We conclude our discussion on the relationship between causal models and causal statements with two key clarifications. First, this study examines the complex interplay between causal statements about events and the causal models that represent connections among propositions (variables and their values). Following Lewis's (1973: 562) perspective, we assume a correspondence between events and propositions: for any event e, there exists a proposition $P(e)$ that holds true in all possible worlds where e occurs. This proposition $P(e)$ affirms the occurrence of e, providing a foundation for articulating the truth conditions of causal statements about events within models that map the relationships between variables.

Second, concerning the linguistic representation of cause and effect, it is noted that in causal statements, the linguistic elements that appears in the relevant sentences act as proxies for the propositions that hold true in specific events. In sentences featuring a CoS verb, the subject DP and the VP represent the CAUSE and EFFECT elements, respectively. For example, in the statement "Sam opened the door," 'Sam' corresponds to condition E in the model (Figure 2), signifying an event where an action leads to the door opening (with variable F assigned a value of 1, denoting the door's open state).

This discussion lays the groundwork for a detailed analysis, in the following subsections, of the semantics of the verb *cause* and CoS verbs, with particular attention to the evaluation of their truth conditions in the context of relevant causal knowledge.

5.2 Differentiating Causal Inferences in Causative Constructions

Let us begin the illustration of our approach with the semantics of overt *cause*. As already claimed by Mackie (1965), and more recently advocated by Nadathur and Lauer (2020), this verb can select as its subject any condition on which the value of the effect causally depends—i.e. any INUS condition. Going back to the example of the electronic door, all of the conditions can be represented as the subject of this construction to describe an door opening, as illustrated in (10), the subscripts in the example refer to the model in Figure 2:

(10) {Sam's pushing the button$_E$/the button$_E$/Sam$_E$/electricity$_D$/the closed circuit$_C$} **caused** {the door to open$_F$}.

Mackie (1965: 247) already clarified that a statement which asserts a singular causal sequence of such a form as "A caused P," makes, implicitly, the following claims:

1. A is at least an INUS condition of P, it is part of a sufficient set B.
2. A was present on the occasion in question.
3. All other members of the sufficient set B were present on the occasion in question.
4. Every other set C that could together bring about P was absent on the occasion in question.

Mackie's claims can be categorized into two distinct types of statements. The first type comprises existential statements that describe the presence or absence of specific conditions at a given moment (2–4). The second type pertains to what we term "causal models," addressing the necessity of particular conditions for an outcome and their inclusion within sufficient sets (1, 3). This type also incorporates knowledge of alternative conditions that could lead to the same result (4). These distinctions are central to understanding the constraints governing the use of the verb *cause*, with (1) specifying the CC-selection criteria for selecting an element from the relevant sufficient set. More broadly, this framework highlights that causal statements report specific events that occurred while adhering to regularities encoded within causal models.

The requirements specified in (1)–(4) are formalized in Formula (11). This formula provides an analytical account of the truth conditions without committing to a lexical definition for the verb *cause.* Instead, it specifies the logical constraints governing CC-selection in the given construction. At this stage, defining the precise lexical entries for these verbs lies beyond the scope of this paper. This limitation reflects ongoing discussion, as highlighted by Nadathur and Bar-Asher Siegal (2024), regarding the extent to which verbs contribute intrinsic meaning versus the role played by aspectual contexts, which impose varying constraints.

Our analysis, consistent with Nadathur and Bar-Asher Siegal (2024), emphasizes the importance of incorporating lexicon-compatible knowledge aligned with the information represented in causal models. Accordingly, we focus on defining the truth conditions for sentences containing inflected predicates, while deliberately postponing a detailed compositional analysis of individual lexical elements.

In the context of example (10), based on the previous discussion, we delineate the components of the truth conditions as follows:

i. Each element within the conditions set {*Sam's pushing the button*$_E$, *the button*$_E$, *Sam*$_E$, *electricity*$_D$, *the closed circuit*$_C$} and the result {*the door to open*$_F$} represents an eventuality. Here, variables are classified as either true or false, or, in our terms, a certain situation holds.

ii. The notation $S(e)$ signifies a situation S (defined as a set of variable-value pairs in Definition 3) that is valid in the context of event e. This implies that the situation S occurs within the event e. A situation may include single or multiple variable-value pairs.

iii. The function SUFF(ICIENT) is applied to two situations, returning 1 if the initial situation S is a sufficient set within model M (see Definition 2 for a Model) for a specific result R, with the second situation represented as a singular pair.

iv. Given that claims (1) and (4) concern causal model condition interactions, we also incorporate second-order logic, allowing for quantification over relations (predicates).

To specify the conditions for a felicitous causal statement in which a particular condition Q corresponds to an element within the causal model, we present the following formula. This formula delineates the truth conditions under which Q can be represented as the subject (the cause) in a causal construction using the verb *cause*:

(11) **Overt *cause***
$\exists Q \exists R \exists e \exists S$: $\text{SUFF}(S, R)^M = 1$ & $(Q \in S^M)$ & $S(e)$ & $\forall O[(O \nsubseteq S)$ & $\text{SUFF}(O, R)^M = 1 \rightarrow \neg O\ (e)]$

The formula in (11) aims to capture the following aspects:

1. It indicates that in the model, situation S forms a sufficient set for situation R. In (10), for example, R represents the pair $< F, 1 >$, indicating that the door is open. A possible subject {Sam's pushing the button$_E$/the button$_E$/Sam$_E$/ electricity$_D$/the closed circuit$_C$} is a representative of the situation represented by Q, which is a member of situation S. Thus, it is a condition characterized as an INUS condition, i.e., a member of a sufficient set. According to Definition 5, all members of this set are also necessary (claim 1).

2. It indicates the presence of an event e where all conditions in the sufficient set are present (claims 2-3).

3. It signifies that no other sufficient sets of conditions were fulfilled in the relevant eventuality (claim 4).[29]

Further nuances will be added to these points in the subsequent discussions. For simplicity, it is presumed that all conditions remain fulfilled at the time of event e. However, it should be recognized that in certain cases, it suffices for the conditions to have occurred prior to the selected cause, initiating the sequence of subsequent conditions.[30]

Shifting our focus to the lexical CoS construction, we contend that its specific constraints governing the selection of a main cause (represented as the subject) give it a different inferential profile from overt *cause*. To see this more clearly, it will be useful to assess the appropriateness of alternative CoS causative descriptions—differing only in their selected subject— within a shared contextual background. This entails identifying the causal descriptions whose parameters align most fittingly with the specific details of the context, based on an underlying causal model. Starting with our model for the operation of an automatic door, as depicted in Figure 2, and the default scenario, in which Sam presses the button and opens the automatic door, we observe a speaker preference for the sentences in (12) over the alternative presented in (13):

(12) {Sam's pushing the button/the button/Sam}$_E$ **opened** the door.

(13) Electricity$_D$ **opened** the door.

But the default preference for (12) over (13) can be reversed if we adjust the contextual background. Imagine applying the same model to a different scenario: Sam pushes down the button but nothing happens, because of a momentary power outage. When power returns, the door opens. Given this alternative scenario, speakers' acceptability judgments shift to prefer (13).

We propose that CoS causatives, in contrast to the overt verb *cause*, are sensitive to a phenomenon that can be perceived as "last straw effect." Specifically, **they require the selection of a condition that completes *a* sufficient set**, ensuring that any remaining necessary conditions within the active causal pathway are also satisfied. In the context of statements of singular causation, which pertain to specific events, the following considerations are essential:

A. Completion involves events taking place in time.
B. Completion also involves a responsiveness to changes in the value of conditions that are linked to events.
C. When assessing a causal statement, it's crucial to differentiate between the transition period when a proposition becomes true and the duration it remains valid. This distinction is vital when attributing an event as the cause. Specifically, the relevant condition's occurrence time marks the point from which we acknowledge the condition's realization. This moment is characterized by a change in the model's associated variables (from 0 to 1 or vice versa), signifying the event's occurrence (refer to Halpern and Pearl 2001, p. 196, AC2; Hobbs, 2005, p. 188).

By incorporating information about time of occurrence causal factors take on a temporal ordering. This allows us to identify the unique **completion event**, the event which "completes" a sufficient

[29]For reasons that will be clarified later, if O contains S, the condition in the situation O can occur as well, as there might be different sizes of sufficient sets according to our definition.

[30]Pragmatic factors can impose additional constraints on the selection of INUS conditions within a specific context. For instance, in a set of causal factors, the least expected one might be selected as the causal subject (for relevant literature on causal selection, see Section 2.1). However, since this paper specifically focuses on CoS causatives, we will not delve deeper into this aspect in the meaning of the verb *cause*.

set, such that following this event (but not before) the values of the set of conditions in the sufficient set entail that the effect occurs.[31]

To clarify, time pertains to the occurrence of events and is not a component of the causal model itself. Instead, it is integral to the truth conditions of specific causal statements that refer to actual causation—that is, to particular events. CC-selection constraints, accordingly, pertain both to the characteristics of conditions within the causal model and to the occurrence of specific conditions in the actual world. The selection pattern of the CoS verb demonstrates its sensitivity to precisely this type of event-related change in condition values within a sufficient set. By default, the subject of the verb is associated with a participant in the **completion event**. In the case of opening an automatic door under normal circumstances, this completion event typically corresponds to the Button condition. As a result, sentence (12) becomes the preferred causal description by default, while sentence (13), which substitutes Electricity as the subject, is deemed unacceptable.

The concept of a completion event is necessary to account for the reversal of judgments, favoring (13) over (12), when the door scenario is altered such that the button is pressed at time t-1, but a power outage delays the door's opening until electricity is restored at time t. In this alternative context, two non-simultaneous event-related conditions are present within the sufficient set (Button and Electricity). This reveals that when the temporal order of events can be retrievable from the context, the felicity conditions of CoS verbs necessitate the selection of the factor associated with the completion event based on the temporal ordering. The power outage scenario establishes a temporal order in which Electricity completes the sufficient set, thus making sentence (13) the most appropriate description.

These findings have been corroborated by Bar-Asher Siegal et al. (2021), whose experiments revealed participants' heightened sensitivity to the completion of a sufficient set when evaluating CoS verbs. The acceptability of causative constructions varied depending on the causal condition referenced. Specifically, CoS verbs were judged appropriate when the causal condition completed the sufficient set. In contrast, when the causal condition was a necessary condition that did not complete the sufficient set, participants found CoS verbs less appropriate than overt causatives.

The constraints on choosing the causative element in a CoS verb are formally detailed in (14). Similar to the earlier discussion on (11), the formulation in (14) exclusively provides an analytical exploration of the truth conditions relevant to CoS verbs. As noted earlier, a further examination of the role of the aspect is essential before we can define comprehensive lexical entries. Therefore, for a felicitous causal statement—where a particular condition Q specifies a condition within the causal model—(14) elucidates the truth conditions for a representative of the condition Q to be selected as the subject (the cause) in a causal structure using a CoS verb. Here, $\tau(e)$ signifies the running time of event e:

(14) **Lexical causative**:
$$\exists Q \exists R \exists e \exists t \exists S{:}\textsc{suff}(S,R)^M = 1 \;\&\; (Q \in S^M) \;\&\; S(e) \;\&\; \tau(e) \subseteq t \;\&\; \forall t' < t \forall e' : \tau(e') \subseteq t' \rightarrow [\neg Q(e')] \;\&\; \forall O[(O \nsubseteq S) \;\&\; \textsc{suff}(O,R)^M = 1 \rightarrow \neg O(e)]$$

The formula in (14) aims to capture the following aspects:

[31] Martin (2018) proposes that the condition selected as the subject of a CoS verb must be "the sufficient cause." This aligns closely with our intuition, as sufficiency viewed as a property of an individual condition implies that the last condition could be perceived as the de facto sufficient cause. However, in Martin's approach, for a condition to be selected as the subject of a CoS verb, it must be an event in which the probability of the occurrence of the effect increases to 1. Without delving into all the details of the differences between our approaches, we will briefly note that her approach encounters problems with cases that we discuss in Section 5.3, where all members of the deterministic chain can be selected as the subject, despite the fact that in most of them, the probability does not change to 1 (it's already 1 before their occurrence). Furthermore, in this approach, once the chain starts to operate, without a clear notion of a causal model (which is not part of her framework), it becomes impossible to distinguish between conditions in the chain and other conditions in the model, as the probability for the occurrence of the effect is 1 in all of them.

Similarly, for her approach, it will be challenging to account for the scenario we discuss in Section 5.6, in which two sufficient sets operate. The first set increases the probability of the occurrence of the result to 1 but is not the actual cause and therefore cannot be presented as the subject.

These challenges for her approach emphasize the advantages of considering the entire model and, especially, considering each condition in the context of its sufficient set.

1. It indicates that in the model, situation S forms a sufficient set for situation R. In (12), for example, R represents the pair $< F, 1 >$, indicating that the door is open. A possible subject {Sam's pushing the button$_E$/the button$_E$/Sam$_E$} is a representative of the situation represented by Q, which is a member of situation S. Thus, it is a condition characterized as an INUS condition, i.e., a member of a sufficient set. According to Definition 5, all members of this set are also necessary.

2. It indicates the presence of an event e where all conditions in the sufficient set are present.

3. It signifies that no other sufficient sets of conditions were fulfilled in the relevant eventuality.

4. It signifies that prior to the time of the event, $Q(e)$ did not hold. Consequently, since Q did not hold prior to τ and all the conditions in the sufficient set are present at τ, the situation Q must be the last one to occur, making event e the completion event, which is when there was a change in the value related to Q (i.e., a state transition).[32]

We can now elucidate the observed asymmetrical entailment pattern between the verb *cause* and CoS verbs in (3) – the starting point of our discussion. When comparing the formulations in (11) and (14), which outline the truth conditions for an explicit *cause* and for the causal component of a CoS verb respectively, it becomes evident that both include the aspects presented in (1-3). However, only (14) encompasses (4), denoting the concept of a completion event. Consequently, (11) is weaker than (14), accounting for the asymmetrical entailment pattern observed. The condition signified by the subject of a CoS verb invariably corresponds to a necessary condition, thus ensuring the truth conditions for the use of the verb *cause*. In contrast, the reverse entailment does not necessarily apply, as *cause* may designate a condition that does not complete a sufficient set.

With the formula in (14) defining the truth conditions for CoS verbs, we now turn to the subsequent subsections, where the various open questions from the literature on these verbs (introduced in Section 4) will be addressed. We will demonstrate how these questions are resolved by our proposed interpretation of the meanings of these constructions.

5.3 Resolving the "directness" connotations of CoS verbs through CC-Selection and sufficient sets

Our analysis offers a key insight into the "directness" connotations associated with CoS verbs, providing a framework to predict the acceptability of sentences such as those in (5), which describe scenarios where additional causal factors intervene between the selected cause and its outcome. We argue that the events represented as subjects in these sentences complete sufficient sets, even when other causal elements subsequently occur.

The apparent challenge posed by the examples in (5)—where multiple points in the causal chain can function as the cause in a CoS verb construction—is addressed by two considerations outlined in the preceding discussion:

1. The process of CC-selection does not assume that only one condition can be selected as the cause (see above 2.1), it determines which ones *can* be selected. Therefore, more than one causal factor can be selected.

2. The CC-selection in the case of CoS verbs is sensitive to the completion of **a sufficient set**, rather than **the completed sufficient set**. The subject must be a representative of the completion event as represented in (14). As emphasized in section 3.2.2, in the case of deterministic

[32] The formula presented in (14) does not preclude scenarios of overdetermination, where two distinct sufficient sets are fulfilled simultaneously. Indeed, there lacks a systematic empirical investigation into whether sentences with CoS verbs are considered true under these circumstances, when the subject represent a cause from only one of these sets. Preliminary, non-systematic surveys among native English speakers yield conflicting outcomes; some find such sentences acceptable, whereas others do not.

causal chains, when each intervening condition is fully determined by a preceding necessary condition, there can be multiple sufficient sets. In fact, there can be as many sufficient sets as there are conditions in the chain, since each of them can be a member of a different sufficient set. Consequently, any stage of this chain, when it occurs, has the potential to complete a sufficient set.

By incorporating these two elements into our proposal, CoS verbs no longer impose restrictions on the selection of a causal factor in the case of deterministic causal chains, as observed in reference (5). When the model guarantees the value of all variables between the designated cause and effect, any non-terminal node in the chain can represent the completion event of a sufficient set. Consequently, each condition in the chain (such as "opening bus lanes" or "accidents increase," in (5)) can be selected as the subject of a CoS verb. The scenario of the automatic door in Figure 2 also allows for variation in the "size" of the sufficient set selected by the CoS verb *open*. (15) demonstrates that the subject can be a participant in either the Button condition or the dependent Circuit condition, since each of these conditions can complete a different sufficient set (as illustrated at the end of Section 3 in Figure 3).

(15) {The pushing of the button$_E$/the closing of the circuit$_C$} opened the door.

In the current proposal, we move away from relying on the intuitive concept of directness to explain the restrictions on the contexts in which CoS verbs can be used. Instead, we elucidate these restrictions through properties that are relative to the model. Thus, employing these verbs requires the completion of *a* sufficient set. The perception of direct causation arises (epiphenomenally) from comparing CoS verbs with explicit *cause* sentences: the stronger selection pattern of the former, which requires a completion event, often excludes more temporally distant conditions, whereas the latter accepts any necessary condition.[33] This gives the impression of a stronger requirement for contiguity in CoS verbs.

5.4 Role of agents

In the context of CoS verbs, other events can occur between the event represented by the cause (the subject) and the resulting effect, as long as they form a completely deterministic causal chain. However, this is not applicable when the intervening events are influenced by intentional agents, as demonstrated by Katz's example involving the sheriff in (16):[34]

[33]In various studies it has been demonstrated that "recency" plays a role in causal selection (Einhorn and Hogarth, 1986; N'gbala and Branscombe, 1995; Henne et al., 2021). However, the fact that any condition in a deterministic chain can equally be selected as the subject of a CoS verb indicates that it is not the time order that matters, but rather being a condition that completes a sufficient set. Additionally, recency is a factor only when other conditions, such as moral responsibility, are met (Henne et al., 2021).

[34]Among the reservations raised by Neeleman and Van de Koot (2012) regarding the direct causation analysis, they present various scenarios in which the Katz effect does not apply. They describe situations where the gunsmith himself can say, "I killed the sheriff," or where others may assert, "The gunsmith killed the sheriff." These scenarios fall into two categories. In one of them, they depict the following scenario:

> "Suppose the gunsmith has held a long-standing grudge against the sheriff and has been pondering how to bring about his death without drawing attention to himself. One day the sheriff brings his gun in for servicing and the gunsmith sees his chance. **Knowing that the sheriff always gets involved in gunfights with unsavoury characters,** he decides to sabotage his gun. Sure enough, the next day there is another gunfight, and as a result of the gunsmith's action, the sheriff's weapon jams at a critical moment and he is gunned down." (p. 9, emphasis is ours)

This case falls under the category of foreseeability, in which a future condition is foreseeable, and we see that earlier conditions can be selected as causes. We will address such cases in Section 5.5 of the paper.

In the case where the gunsmith can say, "I killed the sheriff," the scenario depicts a situation where he is burdened by intense feelings of guilt. We believe that such cases should not be included in the data for a semantic analysis of causation, as they operate in a different manner. People can assert ad-hoc causal relations between two events in almost unlimited ways, and someone can say, "I killed the sheriff" for mystical reasons as well. In all these cases, including the case of the sheriff, one can respond by saying, "Well, you know that in reality it's not you who killed the sheriff, but we

(16) #The gunsmith killed the sheriff. (Repeated from (6))
Context: A sheriff has his six-shooter gun faultily repaired by the local gunsmith. As a result, his weapon jams at a critical moment during a gunfight with a bandit and the sheriff is killed.

Linguistic facts suggest that inexplicit agentive actions are sometimes presupposed to be somewhat deterministic. To see that this is the case, consider the following contrast.

(17) The eclipse ended the performance of Beethoven's 5th.
[*lunar eclipse* > *distracts musicians* > *performance ends*]

(18) While the musicians were performing Beethoven's 5th Symphony, an eclipse coincided, leading to their distraction. After careful deliberation, the conductor made the decision to halt the performance.
?The eclipse ended the concert.

(19) (same context as (17))
The eclipse led to the cancellation of the concert.
[*lunar eclipse* > *conductor stops conducting* > *concert ends*]

In the first case (17), the eclipse is perceived as having a deterministic influence on the musicians, resulting in the concert being unable to proceed. However, in reference (18), the eclipse is understood to trigger a decision on the part of the conductor to cancel the concert. In this scenario, the eclipse leads the conductor to decide, allowing for the possibility that he could have acted differently. As we observe, in such cases, the use of a CoS verb appears to be inappropriate. In contrast, the sentence in reference (19) is formulated using a different causative construction that does not involve a CoS verb, and it is accepted.

The causal model should indicate the relationship of dependency between the eclipse and the cancellation of the concert, as this is the causal knowledge that justifies both (17) and (19). According to our approach to the semantics of causal constructions, it must rely on a model that includes both the eclipse and the conductor's action. However, it is evident that this dependency is not deterministic. There are two possible options: either the model allows for "violations of causal laws" (cf. Schulz (2011)), maintaining a regular causal relation between the eclipse and the conductor's decision, or causal models can also account for non-deterministic connections that allow for fluctuations based on agents' decisions (cf. Nadathur and Lauer (2020)).[35]

Regardless of how SEMs represent indeterminate causal chains, the inclusion of volitional actions plays a critical role in determining what constitutes the completion of sufficient sets. In the examples under consideration, the eclipse itself cannot function as a completion event because all sufficient sets must include the conductor's decision. At the time of the eclipse, no sufficient set is completed that determines the ending of the concert. Consequently, (18) is judged unacceptable, as it fails to select the final condition necessary to complete a sufficient set.

Within the framework of our analysis, the contrast between (17) and (18) demonstrates that the explicit volition of an agent constrains the selection of the cause in CoS constructions, precluding the selection of earlier conditions.

understand why you feel this way." This ability of individuals to make ad-hoc causal claims is related to the broader phenomenon, which is also part of our framework, that causal models are psychological entities. Therefore, speakers can negotiate the structure of these models. As we will discuss in Section 5.5, speakers can construct them in a way that considers a certain action as if it completes the relevant sufficient set. In some cases, the interlocutor may accommodate this perspective, recognizing the specific model for someone experiencing guilt. However, they still have the freedom to insist on the objective perspective and assert that the claim is false.

[35]When causation is indeterminate, causal relevancy (as defined in Definition 5) does not hold. In such cases the dependency, represented in the graph is different, but it is possible to argue that a very similar definition to necessity (Definition 6) still holds (cf. Woodward (2003, 42, 211)), therefore a more flexible definition of causal relevancy would be required. This is, however, beyond the scope of this paper.

5.5 Foreseeablity

Cases in which the satisfaction of one of the necessary conditions can be foreseen appear to violate the requirement for completing a sufficient set. In such situations, a participant in an earlier event can be identified as the cause and become the subject of the clause with a CoS verb. Let's consider again the scenario (discussed earlier as (9)) where Sam is on a train and presses the door button before the train reaches the station. The door only opens when the train stops at the station. In this case, (20) remains acceptable, while (21) sounds less natural, although it may still be considered grammatical by most speakers.[36]

(20) {Sam/the pushing the button/the button} opened the door.

(21) Thee train's arrival at the station opened the door.
Automatic train door with safety delay until the train stops.

Experiments conducted by Bar-Asher Siegal et al. (2021) have indeed confirmed the role of foreseeability in the selection of non-final conditions in the context of CoS constructions. However, given that the completion of a sufficient set is usually a stringent requirement for the acceptability of this construction, our analysis must be adapted, so that a condition can still be considered the "last straw" if, following its occurrence, all other conditions were already foreseeable.

One possible adjustment is to propose the existence of two perspectives from which completion can be considered: an objective perspective and a subjective perspective. From an objective standpoint, the last event that completes a sufficient set is identified (with only time-order being relevant). On the other hand, from a subjective standpoint, the last condition that the agent was unaware of being fulfilled becomes significant (thus, foreseeability is important). This distinction is formally captured in the contrast between (22), which echoes (14), and (23), where the completion event is the last event such that the individual knows about its occurrence. The knowledge is represented as a relation, denoted as K, between a situation S, an individual I, and an event e. This relation signifies the event e in which the individual I become aware of the occurrence of situation S. Thus, we can order the time of the situations according to when their occurrence are known to specifice individuals.[37] The individuals in our case, are the participants in the conversation (i.e. the proposition about occurrence of the foreseeable event is part of the common ground.)

(22) **Lexical causative - the objective take**
$\exists Q \exists R \exists e \exists t \exists S{:}\textsc{suff}(S,R)^M = 1 \;\&\; (Q \in S^M) \;\&\; S(e) \;\&\; \tau(e) \subseteq t \;\&\; \forall t' < t \forall e' : \tau(e') \subseteq t' \rightarrow [\neg Q(e')] \;\&\; \forall O\, [(\text{O} \not\subseteq S) \;\&\; \textsc{suff}(O,R)^M = 1 \rightarrow \neg O\,(e)]$

(23) **Lexical causative - the subjective take**
$\exists Q \exists R \exists e \exists t \exists S{:}\textsc{suff}(S,R)^M = 1 \;\&\; (Q \in S^M) \;\&\; K(S,I,e) \;\&\; \tau(e) \subseteq t \;\&\; \forall t' < t \forall e' : \tau(e') \subseteq t' \rightarrow [\neg K(Q,I,e'] \;\&\; \forall O\, [(\text{O} \not\subseteq S) \;\&\; \textsc{suff}(O,R)^M = 1 \rightarrow \neg O\,(e)]$

An alternative resolution for the challenge of depicting an early event as the subject—provided that a subsequent event (the completion event) is foreseeable—emerges from acknowledging the inherent limitations of causal models. These models are intrinsically incomplete because they invariably omit numerous conditions assumed to be guaranteed. For instance, considering the automatic door scenario, it's evident that various conditions initially excluded from the model play crucial roles: specific positive conditions (such as those related to the electricity supply) and negative conditions (like the lack of obstructions on the door's opposite side). Remarkably, the latter condition becomes

[36]The problem with (21) can be related to the contrast with (20), and to the fact that for some speakers there is always a *preference* for agent as subjects of CoS verbs.

[37]The present analysis does not necessitate the utilization of epistemic logics that employ modal logics, wherein K is understood as a modal concept. In our analysis, K is considered a standard relation, denoting events wherein an individual becomes cognizant of the unfolding of a particular situation. We do not have to take a definitive stance on whether this concept of knowledge requires the semantic apparatus of epistemic logic, typically formulated in terms of possible worlds through Kripke models. This discussion is, however, beyond the scope of the current paper.

pertinent only post-activation of the button, yet is often overlooked in the construction of the relevant model. These overlooked conditions are generally not considered integral to the relevant causal model that licenses the use of a specific causative construction.

Consequently, it is plausible that conditions whose occurrences are taken for granted (including both past and future-foreseeable events) are not regarded as constituents of the relevant sufficient sets. As a result of this, the specific timing of their occurrence becomes inconsequential.

Building on a more principled framework, we propose that selecting the completion event involves identifying the point at which the inevitability of the result's occurrence becomes evident. As a result, it is reasonable to treat foreseeable events as given within these computations, enabling an earlier recognition of the result's necessary occurrence. This perspective aligns with the epistemic approach outlined earlier in our paper (Section 3.1 and throughout), supporting our objective of modeling speakers' causal knowledge within the SEM framework.[38]

5.6 A possible counterexample, and one more definition

What if, in the case of the door equipped with two opening systems (as depicted in Figure 2), it takes 5 seconds from the moment the button is pressed until the circuit closes. Consider the following scenario: John presses the button, but while he waits, Sam turns the handle. In this particular case, who *opened* the door? It is evident that we would attribute the act of opening the door to Sam, while acknowledging that John did not accomplish this task. However, it is noteworthy that John did complete *a* sufficient set, and at the time of its occurrence, no other sufficient set had been completed. As a result, this state of affairs validates the condition represented by (14), implying that the sentence "John opened the door" should be accepted. Conversely, when Sam manually turned the handle, a different sufficient set had already been completed. In accordance with (14), Sam could not be selected as the cause, rendering the sentence "Sam opened the door" unacceptable. Nevertheless, this does not align with the linguistic judgements.

Taking into account the aforementioned considerations, it becomes evident that pressing the button completes *a* sufficient set, as per Definition 6. However, relying solely on this notion proves inadequate for representing a cause in CoS construction. The presence of a 5-second delay mechanism introduces an additional necessary condition that must be satisfied to ensure the desired outcome. Consequently, when another person manually opened the door, this condition remained unfulfilled. Nonetheless, in the case of deterministic chains, as we previously argued, not all necessary conditions need to be fulfilled (and the 5 seconds condition is on a deterministic chain). This realization underscores the necessity of reevaluating the analytical truth conditions for CoS verbs.

A viable solution involves revisiting Mackie's notion of sufficient sets, which cover all necessary conditions for an outcome. Pertinently, we must recall the distinction introduced in Section 3.2.2 between two key concepts: ***a sufficient set*** and ***a completed sufficient set***. The former denotes a collection of conditions that, once known, enable the prediction of a certain result's emergence. Conversely, the latter represents the set which the result materializes only after all its constitutive conditions have actually occurred—this ensemble of factors culminates in the outcome.

To define a completed sufficient set more accurately, we identify all sufficient sets related to the emergence of another condition and then logically ascertain which sets are subsets of others. Within this framework, we pinpoint "the completed sufficient set" specific to a certain interpretation of ψ. To elaborate on this notion, we offer the following refined definition:

Definition 7 (A completed sufficient set): A set of situations which are sufficient sets for

[38]Martin (2018) discusses the notions of *sufficiency ab initio* and *ex post facto* causation, which is also relevant to our analysis. *Sufficiency ab initio* refers to cases where an event, from its initiation, constitutes a sufficient condition for the outcome. In contrast, *ex post facto* causation applies to situations where an event is not initially sufficient but is retrospectively recognized as a cause due to subsequent developments. Importantly, *ex post facto* does not mean that the event, at the time of its occurrence, makes the outcome evident. Instead, it reflects the inclusion of the event as a causal condition in hindsight, based on later unfolding events that factor into the epistemic evaluation. This distinction highlights cases where early conditions, though not immediately sufficient, can retrospectively be identified as causal.

a specific value of ψ, share a completed sufficient set C, if C encompasses all these sets of situations as a superset, and **there is no other sufficient set that contains** C.

The underlying rationale for this definition is as follows. Let us assume that A and B are situations which are sufficient sets for a particular value of ψ, with A being a superset of B ($A \supset B$). In such a scenario, the following conditions must hold:

1. As B is a sufficient set, according to Definition 5, it is causally relevant for the given value of ψ. Hence, even if all other conditions in the model have an undefined value (u), the specific value of ψ can still be determined (as per Definition 4).
2. Given that A is a sufficient set for the same value of ψ, based on Definition 6, every member of A must be necessary for determining the specific value of ψ.
3. The observations made in points 1 and 2 align consistently only in a state of affairs where the values of all members of the the set (A-B), which comprises all members of A that do not belong to B, are deterministically connected to the conditions contained in set B. In other words, while all members of these set are necessary for the occurrence of the result, their values can be "ignored" if they can be deduced from the values of the set B's members.

Taking these considerations into account, we can explore all sets of sufficient sets for a specific value of ψ within a model. Definition 7 captures the fact that by examining which supersets comprehensively encompass all the sufficient sets in the model, we can identify these supersets as the completed sufficient sets for that particular value of ψ in the given model.

Considering this observation, we can conclude that both definitions of **a sufficient set** and **a completed sufficient set** are pertinent for licensing a CoS verb in relation to a specific condition. In a CoS construction, the subject must, on one hand, be part of an event that completes a sufficient set and, at the same time, be a member of **the** (only) completed sufficient set where all the relevant conditions have been satisfied. In the case of the door, pressing the button completes sufficient set 1 in Figure 3, but not all the conditions of its completed sufficient set (set 3) were met. In contrast, turning the handle is part of a completed sufficient set of conditions, wherein all its members were fulfilled.

In light of these observation, we reformulate the criteria that license the cause in the CoS construction. This reformulation, presented in (24) builds upon (14), incorporating the function COMPSUFF takes two situations and returns 1 if the first (Y) is a completed sufficient set in the model for a specific result (R):

(24) **Lexical causative**
$\exists Q \exists R \exists e \exists e' \exists t \exists S \exists Y{:}\textsc{suff}(S,R)^M = 1 \;\&\; \textsc{compsuff}(Y,R)^M = 1 \& Y \supset S \;\&\; (Q \in S^M) \;\&\; S(e) \;\&\; \tau(e) \subseteq t \;\&\; \forall t' < t \forall e'' : \tau(e'') \subseteq t' \rightarrow [\neg Q(e'')] \;\&\; Y(e')\; \forall O\, [(Y \not\subseteq O) \;\&\; \textsc{compsuff}(O,R)^M = 1 \rightarrow \neg O(e')]$

The formulation in (24) aims to elucidate the following aspects:

1. It indicates that in the model, situation S constitutes a sufficient set for situation R. For example, R represents the pair $< F, 1 >$, indicating that the door is open. The subject $\{\text{Sam}_E\}$ is a representative of the situation denoted by Q, which is a member of situation S.

2. It conveys that situation Y in the model establishes a completed sufficient set for situation R, where Y is a superset of S ($Y \supset S$). Consequently, Q is also included in situation Y, characterizing it as an INUS condition according to Mackie—indicative of a completed sufficient set. Per Definitions 6 and 7, all elements within this set are deemed necessary.

3. It asserts the occurrence of an event e wherein all conditions from the sufficient set S are realized.

4. It signifies that prior to the time of the occurrence of event e, condition Q was not realized. Thus, since Q did not realized prior to τ and all conditions in set S materialize at τ, Q must occur last, designating event e as the completion event, marked by a shift in Q's value (i.e., a transition of state).

5. It indicates an eventuality e' where all conditions belonging to the completed sufficient set Y are actualized, with no other completed sufficient sets realized in said eventuality.

Thus, two distinct requirements are identified:

i. The event signified by the subject must be integral to a completion event, entailing that all conditions within a sufficient set have been actualized (4).

ii. It must belong to the sole completed sufficient set in which all conditions were realized (5).

The first requirement implies that the occurrence of the event indicated by the subject **ensured** the outcome's inevitability, rendering it predictable.[39] The second requirement necessitates that this event is included within a constellation of conditions that **directly led to the outcome**. The **predictability** of the result and its **actual production** represent different facets of causal relationships, and both play pivotal roles in the semantics of CoS verbs.

Returning to the scenario mentioned at the beginning of this section, we can now provide an explanation for the linguistic judgments observed. In this scenario, "Sam opened the door" is licensed because Sam participates in an event that fulfilled a condition which is part of the singular completed sufficient set. Turning the handle completes this completed sufficient set, thereby preventing the completion of an alternative completed sufficient set (such as the activation of the circuit), while simultaneously completing a sufficient set, thus being part of a completion event.

6 Conclusions

The common assumption in the linguistic literature is that the category of causative constructions is defined by their shared semantic property of denoting causal relations. However, if we combine this assumption with a unified account of causation, then we should expect the semantics of all causative constructions with similar cause and effect should be the same. This expectation stands in contrast to the fact that different causative constructions have different meanings, as they demonstrate various inferential differences between constructions.

This paper adopts an approach prevalent in linguistic literature to address the seeming contradictions between theoretical assumptions and linguistic data. Within this framework, we encompass all causal relations pertinent to the semantics of causal language under a unified concept of causation. However, we refine this concept by introducing construction-specific requirements. We have advanced this approach by proposing that each linguistic construction is governed by specific constraints—based on its semantics—concerning which conditions from a set in a causal model may represent the selected cause in a causal statement. Thus, our approach to capturing causation unifies through the lens of SEMs, which delineate causal relations in the world. Concurrently, the criteria for selecting causes are defined in terms that relate specifically to the model and to its formal characteristics.

Developing our approach involved modeling causality and constructing a semantic framework to effectively capture the meaning of causative constructions within this model. Following the works of others, we have developed further an application of SEM as a formal model for the semantics of causal statements, enabling us to capture the essential background required for accommodating the

[39] There exists a notable conceptual parallel between the necessity of a completion event for the use of CoS verbs and the semantics outlined by Nadathur and Lauer (2020) for the English verb *make*. According to their analysis, *make* identifies a sufficient causative event that transitions a result from a state of uncertainty to one of certainty within their dynamic framework. Although we have not systematically explored the semantics of *make*, it merits investigation whether there is an intersection between these analyses, potentially revealing a comparative application between CoS verbs and *make*-causatives.

diverse range of causal inferences conveyed through language. Through our investigation, it has become evident that the evaluation of whether a particular condition can be selected to represent the cause is contingent upon its inclusion in a particular sufficient set, rather than being assessed in isolation.

Taking the linguistic perspective, a causal model captures a speaker's knowledge about the relationship between facts in the world, and what is expected to result when certain conditions are fulfilled. With respect to the semantics of causal statements, we argued that this kind of knowledge forms the basis for licensing a speaker's linguistic judgments. Illustrated using two constructions, CoS verbs and overt *cause*, we proposed a hypothesis about the set of constraints imposed by the constructions on the selection of their causes, and how these constraints can be captured by the model. In this way, we shed light on the asymmetry between CoS and overt causatives both in terms of their entailment patterns (3) and in the empirical puzzles related to the notion of causal directness (5)-(7) by elaborating on the restrictions that CoS causatives place on what can be selected as a linguistic subject.

Thus, this approach substantiated the recent claim of Bar-Asher Siegal et al. (2021) that, next to causal selection (the selection of the cause among that various causal factors), there is a causative-construction-selection (CC-selection) as well. In this process the speaker selects, along with the cause, a causative construction which appropriately describes the relation behind the observed course of events. Thus, for a given set of conditions and a given effect, there is a CC-selection that involves the following question: "With respect to each of the relevant condition, can it be encoded as the cause in a statement of a singular instance expressed by a specific causative construction?". Taking this approach, the factors that affect the CC-selection are taken as part of the meaning/truth conditions of linguistic expressions.

It has become evident that CC-selection plays a prominent role in the selection of a statement, next to the process of causal selection. This is primarily due to constraints imposed by the selected linguistic construction. It is recognized that in a given state of affairs certain conditions cannot be represented as causes in specific constructions but can be in others.

The framework of SEM is also a convenient way to represent the fact that causality is a relation between sets of factors, and it allowed us to characterize the selection of the cause with tools provided by this type of model. As part of this analysis, we have also introduced a formalization of a 'sufficient set of conditions' within a model and explained its relevance in selectional parameters. We emphasise the significance of basing sufficiency on sets of conditions which are individually necessary but only sufficient when taken together. We discuss the original motivations for this view from Mackie (1965), and in this paper, these ideas were defined for the first time over the structure of a causal model, and we introduced new supporting evidence from causative expressions in natural languages.

Finally, we have shown that contrastive inference patterns exhibited by causative constructions in English can be precisely captured in relation to SEM with additional characterization of the conditions. Moreover, our analysis was shown to explain longstanding puzzles relating to CoS causative verbs and the intuition behind "direct causation" associated with COS verbs.

As outlined in the introduction, our objective in this paper was to present a comprehensive framework for representing the semantics of causal expressions. The first part of the paper (Sections 2-3) introduced and developed the formal tools within this framework, while the second part (Sections 4-5) provided a test case demonstrating the application of these tools in representing the meaning of two causative constructions. It is our hope that this work will pave the way for future investigations into the broader study of the meanings associated with causal language.

References

Baglini, Rebekah, and Itamar Francez. 2016. The Implications of Managing. *Journal of Semantics* 33. URL `goo.gl/zIhKs1`.

Ballweg, Joachim. 1997. *A model- theoretic account of the concept of causality in natural language*, 223 – 240. Leiden, The Netherlands: Brill. URL `https://brill.com/view/book/9789004373082/BP000008.xml`.

Bar-Asher Siegal, Elitzur, Noa Bassel, and York Hagmayer. 2021. Causal selection - the linguistic take. In *Experiments in Linguistic Meaning 1*.

Bar-Asher Siegal, Elitzur, and Nora Boneh. 2019. Sufficient and Necessary Conditions for a Non-Unified Analysis of Causation. In *Proceedings of the 36th West Coast Conference on Formal Linguistics*, 55–60. Cascadilla Proceedings Project.

Bar-Asher Siegal, Elitzur A., and Nora Boneh. 2020. Causation: From Metaphysics to Semantics and Back. In *Perspectives on Causation: Selected Papers from the Jerusalem 2017 Workshop*, eds. Elitzur A. Bar-Asher Siegal and Nora Boneh, Jerusalem Studies in Philosophy and History of Science, 3–51. Cham: Springer International Publishing. URL https://doi.org/10.1007/978-3-030-34308-8$_1$.

Beavers, John. 2013. Aspectual classes and scales of change. *Linguistics* 51:681–706. URL https://doi.org/10.1515/ling-2013-0024.

Beller, Ari, Erin D. Bennett, and Tobias Gerstenberg. 2020. The language of causation. In *Annual Meeting of the Cognitive Science Society*. URL https://api.semanticscholar.org/CorpusID:219178021.

Benz, Anton. 2006. Partial blocking and associative learning. *Linguistics & Philosophy* 29:587–615.

Ciardelli, Ivano, Linmin Zhang, and Lucas Champollion. 2018. Two switches in the theory of counterfactuals. *Linguistics and Philosophy* 41:577–621. URL https://doi.org/10.1007/s10988-018-9232-4.

Comrie, Bernard. 1981. Aspect and Voice: Some Reflections on Perfect and Passive. *Tense and Aspect* 65–78. URL https://brill.com/view/book/edcoll/9789004373112/BP000005.xml, publisher: Brill Section: Tense and Aspect.

Copley, Bridget, and Phillip Wolff. 2014. Theories of causation should inform linguistic theory and vice versa. In *Causation in grammatical structures*, eds. Bridget Copley and Fabienne Martin, Vol. 52 of *Oxford Studies in Theoretical Linguistics*, 11–57. Oxford: Oxford University Press.

Copley, Bridget, Phillip Wolff, and Jason Shepard. 2015. Force interaction in the expression of causation. *Semantics and Linguistic Theory* 25:433–451. URL https://doi.org/10.3765/salt.v25i0.3482.

Cruse, D. A. 1972. A Note on English Causatives. *Linguistic Inquiry* 3:520–528. URL https://www.jstor.org/stable/4177739, publisher: The MIT Press.

Davidson, Donald. 1967a. *Essays on actions and events*. Oxford: Clarendon Press.

Davidson, Donald. 1967b. The individuation of events. In *Essays on actions and events*. Oxford: Clarendon Press.

De Freitas, Julian, Peter DeScioli, Josh Nemirow, Max Massenkoff, and Steven Pinker. 2017. Kill or die: Moral judgment alters linguistic coding of causality. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 43:1173–1182.

Dixon, R.M.W. 2000. *A typology of causatives: form, syntax and meaning*, 30–83. Cambridge University Press.

Dowty, David R. 1979. *Word meaning and Montague grammar: the semantics of verbs and times in generative semantics and in Montague's PTQ*. Berlin, Heidelberg: Springer.

Einhorn, Hillel J., and Robin M. Hogarth. 1986. Judging probable cause. *Psychological Bulletin* 99:3–19. Place: US Publisher: American Psychological Association.

Escamilla Jr, Ramon Matthew. 2012. *An updated typology of causative constructions: Form-function mappings in Hupa (California Athabaskan), Chungli Ao (Tibeto-Burman) and beyond*. University of California, Berkeley.

Fodor, Jerry A. 1970. Three reasons for not deriving" kill" from" cause to die". *Linguistic Inquiry* 1:429–438.

Hall, Barbara Corey. 1965. Subject and object in modern English. PhD Thesis, Massachusetts Institute of Technology.

Hall, Ned. 2004. Two concepts of causation. *Causation and counterfactuals* 225–276.

Halpern, Joseph Y, and Judea Pearl. 2005. Causes and explanations: A structural-model approach. Part I: Causes. *The British journal for the philosophy of science* 56:843–887. Publisher: Oxford University Press.

Hausman, Daniel M. 1998. *Causal asymmetries*. Cambridge Studies in Probability, Induction and Decision Theory. Cambridge University Press.

Hausman, Daniel Murray. 2005. Causal relata: Tokens, types, or variables? *Erkenntnis (1975-)* 63:33–54. URL http://www.jstor.org/stable/20013346.

Henderson, Robert. 2014. if not for" counterfactuals: negating causality in natural language. In *Proceedings of the 28th West Coast Conference on Formal Linguistics*. Cascadilla Proceedings Project. URL `https://ling.auf.net/lingbuzz/006353`.

Henne, Paul, Aleksandra Kulesza, Karla Perez, and Augustana Houcek. 2021. Counterfactual thinking and recency effects in causal judgment. *Cognition* 212:104708. URL `https://www.sciencedirect.com/science/article/pii/S001002772100127X`.

Hitchcock, C. 2003. Of Humean Bondage. *The British Journal for the Philosophy of Science* 54:1–25. URL

Hitchcock, Christopher. 2020. Communicating Causal Structure. In *Perspectives on Causation: Selected Papers from the Jerusalem 2017 Workshop*, eds. Elitzur A. Bar-Asher Siegal and Nora Boneh, Jerusalem Studies in Philosophy and History of Science, 53–71. Cham: Springer International Publishing. URL `https://doi.org/10.1007/978-3-030-34308-8`$_2$.

Hobbs, Jerry R. 2005. Toward a Useful Concept of Causality for Lexical Semantics. *Journal of Semantics* 22:181–209. URL `https://doi.org/10.1093/jos/ffh024`.

Hoover, Kevin D., and Kevin D. Autor Hoover. 2001. *Causality in macroeconomics*. Cambridge University Press.

Hume, David. 1779. *An enquiry concerning human understanding*, Vol. 2, 3–212. Unknown Publisher. Containing An enquiry concerning human understanding, A dissertation on the passions, An enquiry concerning the principles of morals, and The natural history of religion.

Icard, Thomas F., Jonathan F. Kominsky, and Joshua Knobe. 2017. Normality and actual causal strength. *Cognition* 161:80–93. URL `https://www.sciencedirect.com/science/article/pii/S0010027717300100`.

Jackendoff, Ray. 1972. *Semantic interpretation in generative grammar*. Cambridge, MA: MIT Press.

Katz, Jerrold J. 1970. Interpretative semantics vs. generative semantics. *Foundations of Language* 220–259.

Lakoff, George. 1970. *Irregularity in syntax*. New York: Holt, Rinehart & Winston.

Levin, Beth, and Malka Rappaport-Hovav. 1995. *Unaccusativity: At the syntax-lexical semantics interface*. Cambridge, MA: MIT Press.

Levshina, Natalia. 2016. Finding the best fit for direct and indirect causation: a typological study. *Lingua Posnaniensis* 58:65–82. URL `https://doi.org/10.1515/linpo-2016-0010`.

Lewis, David. 1973. Counterfactuals and comparative possibility. In *Ifs*, 57–85. Springer.

Mackie, John L. 1965. Causes and conditions. *American philosophical quarterly* 2:245–264.

Maienborn, Claudia, and Johanna Herdtfelder. 2017. Eventive versus stative causation: the case of German causal von-modifiers. *Linguistics & Philosophy* 40:279–320. URL `http://search.ebscohost.com/login.aspx?direct=truedb=ufhAN=122941061site=ehost-live`, publisher: Springer Nature.

Mandel, David R. 2003. Judgment dissociation theory: An analysis of differences in causal, counterfactual and covariational reasoning. *Journal of Experimental Psychology: General* 132:419–434. Place: US Publisher: American Psychological Association.

Martin, Fabienne. 2018. Time in probabilistic causation: Direct vs. indirect uses of lexical causative verbs. In *Proceedings of Sinn und Bedeutung 22*.

McCawley, James D. 1978. Conversational implicature and the lexicon. *Pragmatics* 9:245–59.

Mill, John Stuart. 1884. *A system of logic, ratiocinative and inductive: Being a connected view of the principles of evidence and the methods of scientific investigation*, Vol. 1. Longmans, green, and Company.

Nadathur, Prerna, and Elitzur Bar-Asher Siegal. 2024. Modeling progress: Causal models, event types, and the imperfective paradox. In *Proceedings of the 40th West Coast Conference on Formal Linguistics*, 254–264.

Nadathur, Prerna, and Sven Lauer. 2020. Causal necessity, causal sufficiency, and the implications of causative verbs. *Glossa: a journal of general linguistics* 5:49. URL `http://www.glossa-journal.org/articles/10.5334/gjgl.497/`, number: 1 Publisher: Ubiquity Press.

Neeleman, Ad, and Hans Van de Koot. 2012. The linguistic expression of causation. In *The theta system: Argument Structure at the Interface*, ed. Everaert, Martin. Marijana Marelj, and Tal Siloni (eds.), 20–51. Oxford: Oxford University Press.

N'gbala, Ahogni, and Nyla R. Branscombe. 1995. Mental Simulation and Causal Attribution: When Simulating an Event Does Not Affect Fault Assignment. *Journal of Experimental Social Psychology* 31:139–162. URL https://www.sciencedirect.com/science/article/pii/S0022103185710074.

Pearl, Judea. 2000. *Causality: models, reasoning and inference*. Cambridge: Cambridge Univesity Press.

Pinker, Steven. 1989. Learnability and cognition: the acquisition of argument structure. Learning, Development, and Conceptual Change series.

Pylkkänen, Liina. 2008. *Introducing arguments*. Cambridge, MA: The MIT Press.

Rapaport Hovav, Malka, and Beth Levin. 2002. Change of State Verbs: Implications for Theories of Argument Projection. *Annual Meeting of the Berkeley Linguistics Society* 28:269–280. URL http://journals.linguisticsociety.org/proceedings/index.php/BLS/article/view/3842, number: 1.

Rapoport, T.R. 1999. The English middle and agentivity. *Linguistic Inquiry* 30:147–155.

Rett, Jessica. 2020. Manner implicatures and how to spot them. *International Review of Pragmatics* 12:44 – 79. URL https://brill.com/view/journals/irp/12/1/article-p44$_2$.*xml*.

Rose, David, Eric Sievers, and Shaun Nichols. 2021. Cause and burn. *Cognition* 207:104517.

Schulz, Katrin. 2007. *Minimal models in semantics and pragmatics: Free choice, exhaustivity, and conditionals*. AmsterdamInstitute for Logic, Language and Computation.

Schulz, Katrin. 2011. If you'd wiggled A, then B would've changed. *Synthese* 179:239–251.

Shibatani, Masayoshi. 1976. The grammar of causative constructions: A conspectus. In *The grammar of causative constructions*, 1–40. Brill.

Snider, Todd, and Adam Bjorndahl. 2015. Informative counterfactuals. *Semantics and Linguistic Theory* 25:1–17. URL http://journals.linguisticsociety.org/proceedings/index.php/SALT/article/view/25.01, number: 0.

Spirtes, Peter, Clark N. Glymour, Richard Scheines, and David Heckerman. 2000. *Causation, Prediction, and Search*. MIT Press.

Stalnaker, Robert. 1987. Semantics for belief. *Philosophical Topics* 15:177–190.

Stalnaker, Robert C. 1984. *Inquiry*. Cambridge University Press.

Steyvers, Mark, Joshua B. Tenenbaum, Eric-Jan Wagenmakers, and Ben Blum. 2003. Inferring causal networks from observations and interventions. *Cognitive Science* 27:453–489. URL https://onlinelibrary.wiley.com/doi/abs/10.1207/s15516709cog2703$_6$, *_eprint* : *https* : *//onlinelibrary.wiley.com/doi/pdf*/10.1207/*s*15516709*cog*2703_6.

Thomason, Richmond H. 2014. Formal semantics for causal constructions. *Causation in grammatical structures* 58–75. Publisher: Citeseer.

Tooley, Michael. 1987. *Causation: A realist approach*. Oxford: Oxford University Press, Clarendon Press.

VanderWeele, Tyler J., and James M. Robins. 2009. Minimal sufficient causation and directed acyclic graphs. *The Annals of Statistics* 37:1437–1465. URL https://projecteuclid.org/journals/annals-of-statistics/volume-37/issue-3/Minimal-sufficient-causation-an publisher: Institute of Mathematical Statistics.

Waldmann, Michael R., and York Hagmayer. 2013. Causal reasoning. In *The Oxford handbook of cognitive psychology*, Oxford library of psychology, 733–752. New York, NY, US: Oxford University Press.

Wolff, Phillip. 2003. Direct causation in the linguistic coding and individuation of causal events. *Cognition* 88:1–48.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, USA. Google-Books-ID: IFVJAbgySmEC.