Modeling Causal Processes*

Jun Otsuka[†1], Tomoyuki Hayashi[2], Tatsuya Yoshii[3], and Hayato Saigo[4]

[1, 4]Faculty of Social Informatics, Zen University
[2, 3]Department of Philosophy, Kyoto University

Abstract

We offer a category-theoretic representation of the process theory of causality. The new formalism allows process theorists to (i) explicate their explanatory strategies (etiological and constitutive explanations) using the compositional features of string diagrams; (ii) probabilistically evaluate causal effects through the categorical notion of functor; (iii) address the problem of explanatory irrelevance via diagram surgery; and (iv) provide a theoretical explanation for the difference between conjunctive and interactive forks. We also claim that the fundamental building blocks of the process theory—namely processes, interactions, and events—can be modeled using three types of morphisms. Overall, categorical modeling demonstrates that the philosophical theory of process causality possesses scientific rigor and expressive power comparable to those of its event-based counterparts, such as causal Bayes nets.

1 Introduction

Philosophical positions on causation are primarily divided into two main categories: event-based and process-based. Event-based theories conceive of causation as a specific type of relationship between events; for instance, the intake of aspirin causing relief from a headache. While these theories vary concerning the nature of the relationship that qualifies one event as the cause and another as the effect—be it a constellation of logical conditions, probabilistic regularities, counterfactual or functional dependence, or another factor—they concur on the ontological premise that causal relata are events. Process-based theories, in contrast, take processes, objects, or entities as the fundamental elements of causation. According to this perspective, causation is a specific type of activity or behavior of a cause-thing that affects an effect-thing. Again, the nature of this activity varies among theorists: some perceive it as a physical interaction, others as the expression of an underlying power or tendency of the cause, and so on. Nonetheless, these theories share a common understanding that a causal relationship is something that obtains among *things*, not events.

Glymour (2004) distinguishes between two argumentation styles in the philosophical literature: the Socratic and the Euclidean. The Socratic tradition aims to

*Accepted for publication in *Synthese*.
†Corresponding author: jun_otsuka@zen.ac.jp

clarify what constitutes causality through conceptual analysis, case studies, counterexamples, and the search for necessary and sufficient conditions. Its fundamental question is definitional: 'Causality is about X. Then what is X?' Event-based theories identify X as a specific relationship between events or properties, and then proceed to investigate the nature of this relationship. On the other hand, the process-based view, upon taking X to be a specific type of process, sets out to determine the characteristics of genuine causal processes. Is a shadow a causal process? If not, why? Process theorists in the Socratic mode have sought to answer such questions by providing a conceptual or empirical criterion of causal processes based on mark transmission (Salmon, 1984), conserved quantities (Dowe, 2000), or other properties.

The Euclidean tradition, on the other hand, seeks to model causality in the "form of formal or informal axiomatic systems (Glymour, 2004, p. 779)" without much concern for what causality really consists of. Glymour takes causal inference frameworks developed in the late 20th century, such as causal Bayes nets and Rubin's counterfactual model, as paradigmatic examples of this category. By using directed acyclic graphs (DAGs) or potential outcomes, these techniques provide quantitative methods to identify causal relationships and evaluate intervention effects on the basis of observational data. By and large, these formal methods stand on the event-based view. Random variables are a probability-theoretic expression of events, and edges in a causal DAG assert that events described by one variable are in specific—probabilistic and/or counterfactual—relationships with events described by another variable.

The success of these formal approaches has undoubtedly rendered great credibility to the event-based conception of causality. In contrast, to the best of our knowledge, no similar development of the Euclidean kind has been made in the process camp. This lack of formalism has limited the process theory from advancing into a rigorous and systematic scientific methodology beyond a philosophical speculation that, even if plausible, remains largely intuitive. The main goal of this article is to fill this lacuna, by resorting to the category-theoretic approach to causal modeling developed lately (e.g. Fong, 2013; Coecke and Kissinger, 2017; Jacobs et al., 2019; Fritz and Klingler, 2022; Lorenz and Tull, 2023). In particular, we claim that its graphical representation, *string diagrams*, offers formal definitions of processes and interactions, much like directed acyclic graphs (DAG) serve as formal representations of events and causal relationships in causal Bayes nets.

The benefits of the new formalism are several-fold. First, the compositional features of string diagrams provide a natural way to formulate the two principal explanatory strategies of process theorists, namely etiological and constitutive explanations. Second, the categorical notion of *functor* enables the integration of process theory with probabilistic analysis. Third, by incorporating the intervention calculus, the categorical modeling resolves the issue of explanatory irrelevance, which has been a significant problem for process theory. Fourth, it also captures the crucial difference between what Salmon called conjunctive and interactive forks, which are indistinguishable in conventional DAGs, and explains why the principle of common cause (PCC) fails in the latter. Finally, the categorical framework offers formal definitions of causal processes, interactions, and events, thereby clarifying the ontological differences between event-based and process-based theories.

We begin with a brief sketch of the philosophical theories of process causality in Section 2. Section 3 introduces string diagrams—the category-theoretic apparatus we use to model causal processes and interactions. Section 4 connects the diagrammatic representation of causal structures to a probabilistic analysis via a functor.

Section 5 addresses the issue of explanatory irrelevance through interventions on string diagrams. In Section 6, we propose a categorical modeling of conjunctive and interactive forks, and explain why the common cause in the latter fails to screen off its effects from each other. Section 7 turns to the ontological thesis of process theory and defines processes, interactions, and events in categorical terms of morphisms. Section 8 concludes.

2 Philosophical theories of process causation

In his seminal book, Wesley Salmon (1984) set out to replace what he identifies as the "standard picture of causality ... that we have two (or more) distinct events that bear some sort of cause-effect relations to one another," with an alternative view that "take[s] processes rather than events as basic entities (p. 137)." While events are localized at points in space and time, processes have a greater time duration and are expressed by (bunches of) world lines. A baseball hit by a bat and the pulse of light coming from a distant star are examples of processes. By contrast, a light spot projected onto a wall, though it persists in time and is represented geometrically by a world line, does not qualify as a causal process: it is a *pseudo-process.* The difference between causal and pseudo processes resides in the possibility of *causal interaction.* An interaction between two or more genuine causal processes produces some lasting "mark" or modification in their structure, as the ball hit by the bat gains momentum and perhaps gets a bit of scratch on the surface. This is not the case for pseudo-processes: superimposition of two light spots may momentarily change their color, but this change disappears as soon as they separate. Salmon has us see the world as a bunch of causal processes interacting with each other, and he argues that this worldview is better suited to analyzing causality than the event-based ontology.

Salmon's work has laid the foundation for the subsequent development of mechanistic philosophy and has served as a springboard for a variety of its later offshoots. Phil Dowe (2000) proposed replacing Salmon's criterion, which essentially relies on the counterfactual concept of "marking," with the physical concept of a conserved quantity. In his view, a causal process is the world line of an object possessing a conserved quantity, while a causal interaction is an intersection of world lines involving the exchange of a conserved quantity. What distinguishes a genuine causal process such as a ball from a moving light spot is that the former, unlike the latter, possesses certain quantities such as momentum that are transferred to another object upon interaction. As this analysis suggests, Dowe's primary focus (and to a large extent, Salmon's) is on physical science. In contrast, *New Mechanist* philosophers have developed a similar causal theory focused on life sciences, analyzing biological mechanisms as orchestrated activities and interactions of constituent factors (e.g. Glennan, 1996; Machamer et al., 2000; Bechtel and Abrahamsen, 2005). For instance, photosynthesis is mechanistically explained by identifying how chemical substances and biological components—such as water, carbon dioxide, and chlorophyll—interact to produce oxygen and glucose. Despite the difference in focus and terminology, this approach shares Salmon's view that the analysis of causality should rest upon processes (entities, components) and their interactions (activities), rather than on events and their relationships.

These philosophical analyses attempt to identify the essence of causation and mechanisms in processual language. Dowe (2000), for instance, explicitly frames his goal as an "empirical analysis" that "seeks to establish what causation in fact *is* in the actual world. (p. 3, emphasis in the original)." Mark-transmittable processes, conserved quantities, and entities engaging in activities can be thought of as answers to this what-question. In the shadow of these ontological inquiries, yet almost nothing has been said about how the causal processes and interactions understood as such can be *modeled*, preferably within a certain mathematical framework. We believe that this lack of formal modeling underlies several significant challenges to the process-based view, one of which is the problem of explanatory relevance. As Woodward (2003, p. 357) notes, the mere existence of a causal process "tells us nothing about which features of the process are causally or explanatorily relevant to the outcome." A swinging bat transmits various marks and conserved quantities to a ball, but the theory is silent on which of these are causally relevant to the ball's subsequent motion. A satisfactory solution to this problem requires a suitable modeling framework that supports the analysis of contrast classes and their relevance to the intended outcome. Another related criticism concerns the philosophical purport of the theory: that is, besides providing a metaphysical description of causal relationships, it is not clear how it "can be made to do real philosophical work (Hitchcock, 2004)." For instance, one of Salmon's main motivations for process theory was the distinction between two kinds of fork structures, conjunctive and interactive (which we discuss later in detail). However, Salmon does not explain how the ontological difference between the two forks gives rise to their respective probabilistic patterns, namely that the common cause screens off its effects from each other only in the conjunctive fork but not in the interactive fork. Without such an explanation, process theory remains on a par with event-based theory, and this alleged anomaly cannot be used to motivate his alternative framework.

The present paper takes on this modeling challenge with the aid of the recently developed category-theoretic approach to causal modeling.[1] Our discussions focus more on philosophical implications than on theoretical details; readers interested in the theory itself are referred to Jacobs et al. (2019), Fritz (2020), and Lorenz and Tull (2023). Through these discussions, we argue that the categorical framework provides theoretical substance to the philosophical theory of process causality as well as the means to address its problems.

3 Modeling Process Causality with String Diagrams

The first step in modeling process causality is to represent processes and interactions using *string diagrams*, as employed in the category-theoretic approach to causal modeling (Coecke and Kissinger, 2017; Jacobs et al., 2019; Fritz and Klingler, 2022; Lorenz and Tull, 2023). String diagrams are composed of *wires* and *boxes*, which symbolize causal processes and interactions, respectively.

Wires are drawn as vertical lines, with time flowing from bottom to top (or from left to right when drawn horizontally). We let wires represent what process-theorists call causal processes or physical objects, such as a moving ball or a pulse of light. We stress, however, that the formalism itself is noncommittal to the ontological nature of causal processes. Our aim in this paper is to provide a model for what theorists regard as causal processes, leaving the criteria for distinguishing processes from non-processes to empirical investigation. The substrate neutrality of the formalism

[1] The categorical framework is sometimes referred to as the "process theory (e.g. Coecke and Kissinger, 2017)." To avoid confusion, we use "process theory" exclusively to refer to the philosophical view of causation.

does not undermine our modeling purpose, just as constructing Bayesian networks does not require resolving what can or cannot be represented by random variables.

Another important difference from the previous philosophical approaches is that we allow wires to take on various states or values, as will be elaborated upon in Section 4. In this sense, a wire in a string diagram represents a *type* of causal process, rather than a specific token or concrete instance. The 'ball' wire does not represent a single trajectory of a particular ball; rather, it may correspond to different trajectories in various directions and under distinct conditions. The type interpretation is a natural requirement stemming from the fact that a string diagram is intended not as a mere description of individual phenomena, but as a model that accommodates predictions and counterfactual reasoning.

Next, interactions between wires/processes are depicted by boxes. Each box specifies the type of processes it accepts as inputs and produces as outputs. The following boxes illustrate three different types of interactions, which Salmon (1984, p. 203) referred to as x-, λ-, and y-type interactions, respectively.

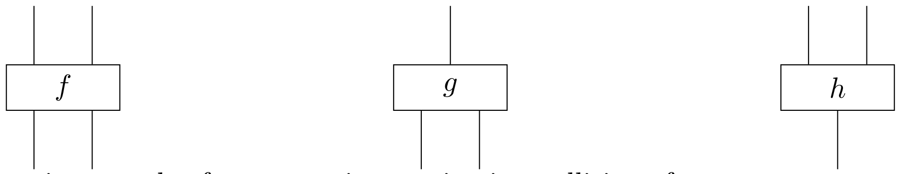

A paradigmatic example of an x-type interaction is a collision of two processes, such as two billiard balls. In a λ-type interaction, two processes converge to produce a single outcome (e.g., the synthesis of water from hydrogen and oxygen). Conversely, a y-type interaction illustrates cases where a single process 'splits' into multiple processes (e.g., mitosis and atomic decay). Interactions involving more than two processes can also be considered and are represented by boxes with multiple input and output wires.

In modeling an interaction we may decide to focus on one process and ignore others. Imagine a machine that stamps a quality-inspected mark on balls in a ball manufacturing factory. This is an x-type interaction between the ball process and the stamp machine process. But if we are interested in only the balls and not the machine, it is reasonable to model the production line with a box having just one input and one output, which represent the balls that enter and exit the machine, respectively:

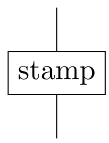

.

Let us call this I-type box. An I-type box may indicate either a spontaneous change in one process, or a change arising from an interaction with other processes that are abstracted away as background conditions for the purpose of modeling.

One particularly important variety of I-type box is *identity*, which leaves the process as it is—i.e., it does not change anything. Identity is depicted by an empty (dashed) box and can, in fact, be identified with the wire itself, as shown below:

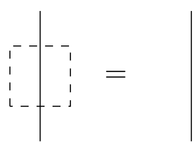

.

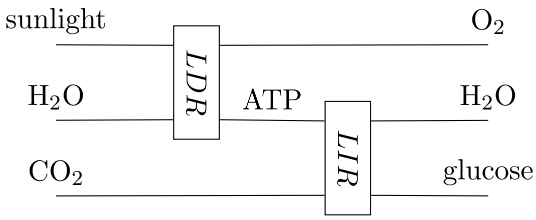

Figure 1: A string diagram representing photosynthesis, which takes sunlight, water, and carbon dioxide to produce oxygen, water, and glucose, through chemical reactions represented by the two boxes. Here the causal flow is from left to right.

This means that any process can be understood as a “null interaction” or the absence of interaction with any other processes (although we will later introduce an alternative interpretation of a wire in Section 7).

It is also possible to abstract away the origins of processes. Consider balls being thrown by a pitching machine—an x-type interaction involving both the balls and the machine. However, one might be interested only in the balls that are ejected. We can model this interaction using a special unit that has no inputs, termed *states*. A state is denoted by a downward triangle, and represents the commencement of a process without detailing its origin. For instance, the following diagram depicts a collision between two balls, where the balls are initially set in motion by their respective states.

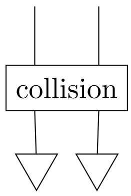

Wires and boxes (including their special variants like identities and states) are the basic building blocks of a string diagram. The second step of categorical causal modeling is to organize these elements into a coherent structure that represents the phenomena under study. Let us proceed with an example. Photosynthesis is a complex mechanism consisting of multi-stage chemical interactions among sunlight, water, carbon dioxide, etc. In the first stage, known as the light-dependent reaction, the thylakoids in chloroplasts produce chemical energy (ATP and NADPH) using sunlight and water. The produced energy is then used to convert carbon dioxide into glucose in the second stage, the light-independent reaction. In the terminology of process theory, sunlight and chemical substances are causal processes, while the two reactions are their interactions. By connecting these processes (wires) and interactions (boxes) in matching types and orders, the entire photosynthesis mechanism can be represented as shown in Fig. 1.

The diagram in Fig. 1 is constructed by connecting wires and boxes in two ways: sequentially or in parallel. The diagram on the left below shows the sequential connection, where the output of box f is fed into subsequent box g. Composition means that the stacked boxes can be seen as one stretch of a mechanism $g \circ f$ that takes A as the input and outputs C, as indicated by the dotted box. On the other hand, the right diagram illustrates parallel processing, where two processes undergo f and h independently. Here again, we can see these two parallel mechanisms as one combined mechanism, $f \otimes h$, that takes the joint input $A \otimes X$ and yields the joint output $B \otimes Y$.

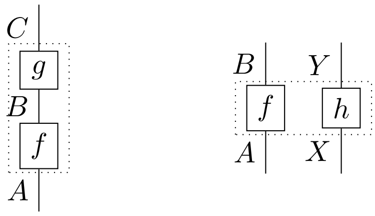

It is assumed that the order in which these two compositional operations are applied does not matter: hence in the following diagram, whether one horizontally combines f's and g's and then stacks them vertically (as on the left), or forms two sequential lines and then puts them side-by-side (as in the middle), the result is the same large box on the right with the combined input $A \otimes X$ and output $C \otimes Z$.

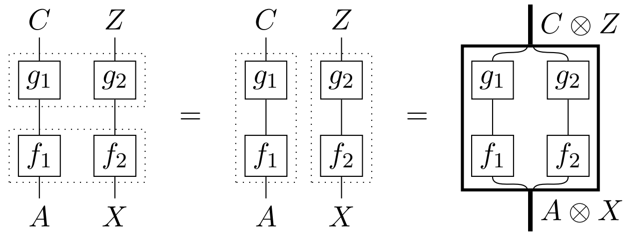

(1)

These compositional features are notable characteristics of the *monoidal category.* In categorical terminology, wires and boxes are called *objects* and *morphisms* of a category, respectively. It is customary to denote objects with capital letters $A, B, C, \dots,$ while morphisms with arrows $f : A \to B, g : B \to C$, and so on. Two morphisms having a matching codomain/domain can be combined to yield a morphism from the domain of the first morphism to the codomain of the second morphism, $g \circ f : A \to C$, which amounts to the vertical composition discussed above. The horizontal composition, on the other hand, exploits the *monoidal* property. A monoidal category is a category equipped with a binary associative operation $\otimes$ that creates, (i) for any pair of objects A and B, the product object $A \otimes B$; and (ii) for any pair of morphisms $f_1 : A \to B$ and $f_2 : X \to Y$, the product morphism $f_1 \otimes f_2 : A \otimes X \to B \otimes Y$. This operation underlies the horizontal or parallel combination. The consistency of the two operations as illustrated in equation (1) is then expressed as their interchangeability,

$$(g_1 \otimes g_2) \circ (f_1 \otimes f_2) = (g_1 \circ f_1) \otimes (g_2 \circ f_2), \tag{2}$$

which is ensured by the axioms of the monoidal category. In this way, string diagrams form a category known as an *affine CD-category* (Jacobs et al., 2019) or a *Markov category* (Fritz, 2020), special types of symmetric monoidal categories augmented with additional structure tailored to represent causal processes.[2]

The categorical modeling accommodates the explanatory strategy of mechanistic philosophy, according to which phenomena are explained by being shown to "fit into a causal nexus (Salmon, 1984, p. 19)" in two different ways. The first is what Salmon calls *etiological explanation*, which places "the explanandum in a causal network consisting of relevant causal interactions (p. 267)." This entails

[2]The symmetric property means that any two parallel processes X, Y may be swapped, so that there is a natural isomorphism $X \otimes Y \xrightarrow{\sim} Y \otimes X$. The category of string diagrams is also equipped with additional morphisms called discard and copier, the latter of which will be discussed in detail in Section 6.

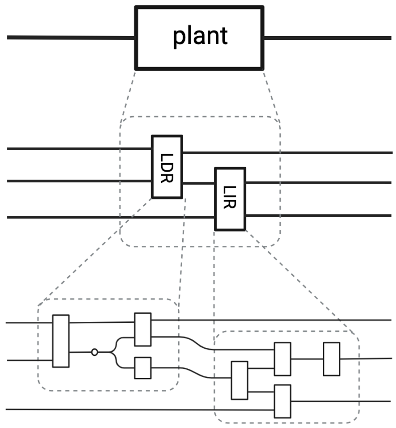

Figure 2: A diagrammatic representation of a constitutive explanation of photosynthesis. A single wire or box in a higher-level description can be considered as consisting of multiple wires or a diagram in a lower-level description. The compositional properties of a monoidal category ensure that such composition is consistently and uniquely defined.

identifying a causal structure that elucidates how various factors interact to yield the phenomenon being explained. In a similar vein, Machamer et al. (2000, p. 12) identify a causal explanation with a portrayal of a mechanism that "describes the relevant entities, properties, and activities that link them together, showing how the actions at one stage affect and effect those at successive stages." In the categorical framework, this amounts to building a string diagram for the phenomena to be explained by combining relevant wires and boxes, as we just did for the photosynthesis mechanism above in Fig. 1.

On the other hand, *constitutive explanation* seeks to elucidate how "the fact-to-be-explained is constituted by underlying causal mechanisms (Salmon, 1984, p. 270)." This is a sort of reductive explanation of the functioning of the higher-level process or interactions in terms of lower-level ones, which is a standard explanation strategy in biomedical sciences (Craver, 2001; Bechtel, 2006). The photosynthesis mechanism sketched above treated the two main reaction systems as black boxes, but each can be further analyzed in terms of its constituent factors; e.g., the light-independent reaction is decomposed into various enzymes (such as RuBisCO and phosphoglycerate kinase) and other chemical compounds, with their interactions being explained by more detailed physicochemical reactions.

A monoidal category offers a natural framework for representing such reductive explanations. We previously noted that wires and boxes can be combined to form a joint wire or box (1). Conversely, a given element in a string diagram can be decomposed into a lower-level string diagram that comprises fine-grained parts. For instance, at the coarsest level, photosynthesis can be depicted by a single box representing a plant that converts one set of chemicals into another (Fig. 2, top). The photosynthesis diagram in Fig. 1 can then be seen as a high-level decomposition of this single box (middle). Similarly, each component of this diagram can be further represented as a composed string diagram involving micro-level wires and boxes, collectively elucidating the mechanisms underlying the light dependent and independent reactions (bottom). In this way, the compositional nature of a monoidal category allows us to formulate the constitutive explanatory strategy, as envisioned by Craver (2001, p. 66).

In summary, string diagrams provide a formal apparatus to capture the ontological intuition of process theorists that causation is more effectively understood in terms of processes and their interactions, rather than events and their relationships. This formalism accommodates the two primary explanatory goals of mechanistic philosophy: etiological and constitutive explanations, owing to the inherent compositional nature of monoidal categories. However, the exposition thus far remains merely a pictorial *description* of process causation. Comprehensive *modeling* of a causal structure should be complemented by a quantitative assessment, one that evaluates the extent and probability with which causes influence their effects. This is the point at which the categorical formulation demonstrates its usefulness, as we discuss next.

4 Probabilistic Interpretation

String diagrams for process-based theory are analogous to DAGs for event-based theory in that they both encode qualitative information about which items cause which. Comprehensive Bayesian network modeling involves pairing a DAG with a probability distribution that satisfies conditions such as the Markov condition. Similarly, a string diagram must be given a probabilistic interpretation for the quantitative analysis of process causation. This is achieved through a *functor*, a systematic mapping from one category to another. In the present context, we consider a functor F from the category C of string diagrams to that of finite sets and stochastic matrices, which effectively assigns values to each wire and a conditional distribution to each box.

We illustrate the functorial assignment with the photosynthesis mechanism described in Fig. 1 (the following depiction is merely for illustration purposes, with no claim for truth). The first part of the interpretation is to associate each process (wire) with the set of possible values it can take. For instance, $F(sunlight) = \{+, -\}$ indicates that *sunlight* can be either abundant (+) or scarce (−). For the sake of simplicity, let us assume that all the other five types of wire in the photosynthesis diagram are also binary, so that $F(H_2O) = F(CO_2) = F(O_2) = F(ATP) = F(glucose) = \{+, -\}$.

Next, we will substantiate the boxes by specifying how the interactions they represent affect the involved processes. This is achieved by assigning to each box a conditional distribution that determines the probability of its output values given its input values. The photosynthesis diagram contains two boxes, LDR and LIR, each having two inputs and two outputs. Let us assume that their operations are described by the following conditional distributions:

LDR

$sunlight$	+		−	
H_2O	+	−	+	−
O_2^+, ATP^+	.7	.2	.1	0
O_2^+, ATP^-	.1	.1	.2	.1
O_2^-, ATP^+	0	.3	.1	0
O_2^-, ATP^-	.2	.4	.6	.9

LIR

ATP	+		−	
CO_2	+	−	+	−
H_2O^+, glu^+	.8	.5	.2	0
H_2O^+, glu^-	.1	.1	.2	.1
H_2O^-, glu^+	.1	.3	.1	0
H_2O^-, glu^-	0	.1	.5	.9

Table 1: Conditional distributions assigned to the LDR (left) and LIR (right) by functor F.

Each cell in the table indicates the conditional probability of a pair of output

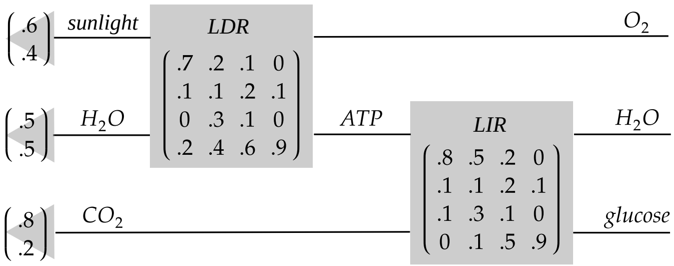

Figure 3: A hypothetical example of the probabilistic interpretation of the photosynthesis diagram in Fig. 1.

values, as listed in the leftmost column, given the combination of input process values shown in the top rows. For instance, the top-left cell of the LDR table indicates that in the presence of abundant sunlight and water, the LDR produces a good amount of O_2 and ATP with a probability of 0.7. If a box has i inputs, each with $n_1, n_2, \ldots, n_i$ values, the total number of input combinations is $N = \prod_{k=1}^{i} n_k$; likewise if it has j outputs, each with $m_1, m_2, \ldots, m_j$ values, the total number of output combinations is $M = \prod_{l=1}^{j} m_l$. The box is then assigned an $M \times N$ matrix whose entries are probability masses. Such a matrix is called a *stochastic matrix*. In general, the interpretation map F assigns a stochastic matrix of corresponding dimensions to each box f.

The same procedure applies to states, represented by downward triangles that signify the commencement of processes. In the probabilistic interpretation, a state is assigned a marginal distribution of its output. This is expressed as an $M \times 1$ stochastic matrix (i.e., an M-vector), where M represents the number of values of the output wire. In the present case, F maps the states $s_{sunlight}$, s_{H_2O}, and s_{CO_2} to 2-vectors that represent marginal distributions over the abundance of the respective elements.

Recall that the identity in a string diagram is a special type of box that performs no operations and can thus be equated with the wire itself. For a process with M values, the identity is interpreted by an M-dimensional identity matrix I_M, which has ones on its diagonal and zeros elsewhere. Since multiplying an identity matrix does not change a vector (and thus a vector representation of a probability distribution), i.e., $I \cdot v = v$ for any vector v of the matching dimension, this effectively illustrates the null process that leaves its input unchanged.

Fig. 3 shows a complete photosynthesis mechanism interpreted with functor F. The interpreted diagram allows one to calculate how "causal processes transmit probability distributions (Salmon, 1984, p. 262)" via matrix calculus. For instance, the joint distribution $p(O_2, ATP)$ is given by $F(LDR) \cdot F(s_{sunlight}) \otimes F(s_{H_2O})$, where $\cdot$ is matrix multiplication and $\otimes$ is the Kronecker product. To obtain this, we first calculate the joint distribution of *sunlight* and H_2O by $F(s_{sunlight}) \otimes F(s_{H_2O}) = (.6, .4)^T \otimes (.5, .5)^T = (.3, .2, .3, .2)^T$. Then multiply this vector with $F(LDR)$ given by the left Table 4 above, which yields $(.28, .13, .09, .5)^T$ whose values respectively give the probabilities of (O_2^+, ATP^+), (O_2^+, ATP^-), (O_2^-, ATP^+), and (O_2^-, ATP^-). The probability outcome of the LIR can be obtained in a similar fashion.[3]

The finite sets and stochastic matrices used here to interpret a string diagram

[3]For detailed explanation and more sophisticated reasoning such as Bayesian conditionalization, see Jacobs, B. (unpublished). *Structural probabilistic reasoning*, `https://www.cs.ru.nl/B.Jacobs/PAPERS`.

form a category called FinStoch. The probabilistic assignment introduced above thus maps objects (wires) and morphisms (boxes) of category C of string diagrams to those of FinStoch. Such a systematic mapping F from one category C to another is called *functor*, if it preserves (i) identity, so that $F(Id_A) = Id_{F(A)}$ for any object A of C, and (ii) the composition of morphisms, so that $F(f \circ g) = F(f) \circ F(g)$ for any morphisms f, g of C.[4] In the present context, (i) means that the identity wire (empty box) is mapped to an identity matrix, while (ii) is guaranteed by the rule of matrix multiplication. Thus in the categorical approach, a pair of category C of string diagrams and a functor from it to FinStoch defines a causal model.

5 Intervention and Explanatory Relevance

One salient feature of the Bayesian network is its ability to estimate hypothetical interventions via graph manipulation or *do* calculus (Spirtes et al., 1993; Pearl, 2000). In the categorical framework, an intervention is implemented by a diagram surgery that replaces a box with a new state that represents a post-intervention distribution (Jacobs et al., 2019). For instance, suppose that a scientist is interested in how glucose production is affected by watering the plant. This can be estimated by replacing s_{H_2O} with a new state s'_{H_2O}, whose probability interpretation $F(s'_{H_2O})$ assigns a probability of one to H_2O^+ and zero to H_2O^-. The post-intervention probability distribution for the remaining parts can be calculated on the basis of this state by applying the matrix calculation outlined in the previous section.

The intervention calculus provides a means to address one of the most significant challenges raised against process theory, namely the problem of explanatory irrelevance (Woodward, 2003; Hitchcock, 2004). The issue is that process theory cannot distinguish factors that are relevant and responsible for a given effect from those that are not. Consider Salmon's example of John Jones, who regularly takes birth control pills and fails to become pregnant. Although John's inability to conceive is due to biological sex rather than medication, birth control pills still represent genuine causal processes that interact with his body when ingested. More broadly, any event can be understood as the culmination of myriad processes that travel within the timelike region of the light cone, though most of these processes have no influence whatsoever on the occurrence of the event. Therefore, identifying some or even all causal processes involved in a particular event still falls short in explaining why that event occurred, which is arguably a primary motivation for considering causal relationships.

The causal irrelevance can be understood as insensitivity to an intervention. Let us illustrate this with the above episode of John Jones, modeled in Fig 4. This simple model stipulates that pregnancy Z is determined by the interaction between two causal processes, regular intake of birth control pills X and the reproductive organ Y. Assume they are all binary, with $F(X) = \{pill^+, pill^-\}, F(Y) = \{male, female\}, F(Z) = \{pregnant^+, pregnant^-\}$, where + and − indicate the presence and absence, respectively. Then, being a male is represented by the state s_Y that assigns a unit probability to *male*. Under this condition, any intervention on X, represented by the replacement of the original state s_X with a new state s'_X, does not affect the probability of pregnancy, for surely $p(pregnant^+|pill^+, male) = p(pregnant^+|pill^-, male) = 0$. In contrast, a similar intervention on X should change the outcome if the subject is female, for arguably $p(pregnant^+|pill^+, female) \neq$

[4]Moreover, in the present context, a functor must be a monoidal functor that preserves the monoidal structure.

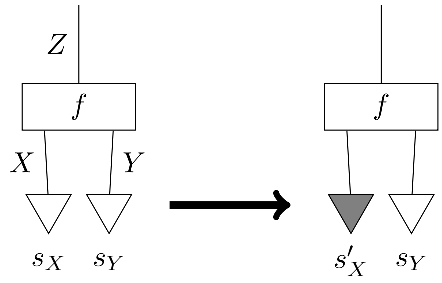

Figure 4: Intervention on X that replaces the original state s_X with the new state s'_X.

$p(pregnant^+|pill^-, female)$. This difference between intervention outcomes bears out our intuition that taking pills is irrelevant to John's barrenness.

In general, a causal process X is *explanatory irrelevant* to an outcome Z under a state $s_\mathbf{Y}$ of other processes $\mathbf{Y}$ if any intervention on X does not affect the marginal distribution of Z under $s_\mathbf{Y}$. In the *do*-calculus formalism, this amounts to $p(Z|do(x), \mathbf{Y}) = p(Z|do(x'), \mathbf{Y})$ for any value $x \neq x'$ of X. By incorporating the intervention calculus by way of the diagram surgery, the categorical framework harnesses the process theory with a means to handle explanatory irrelevance, in the same way as the causal DAG framework.

6 Interactive vs. Conjunctive Forks

A major impetus for Salmon's process causal theory comes from the problem of common causes (Salmon, 1984, ch. 6). Reichenbach's statistical explication of the principle of common cause (PCC) suggests that two probabilistically dependent events A and B become independent (or screened-off from each other) if conditioned on their common cause C so that $P(A, B|C) = P(A|C)P(B|C)$, provided that there are no other causal connections. Salmon's examples include two essays submitted by students who plagiarized from the same paper in a fraternity file, and dining companions who experience severe stomach discomfort after eating the same meal. In both cases, the unlikely correlations—similar word usage and sickness in otherwise healthy guests—would certainly vanish if we consider their common causes: the source paper and the tainted food.

However, Salmon (1984, pp. 168-174) indicates that in some cases a common cause fails to screen-off a statistical dependence. Imagine a billiard player attempting to pocket an 8-ball by striking the cue ball. However, given their arrangement, if the 8-ball sinks into the pocket (A), the cue ball will inevitably end up in the opposite pocket (B) as well. In this scenario, A and B remain dependent even when conditioned on their common cause, namely their collision (C).

From this, Salmon concludes that there are two distinct common-cause structures: the *conjunctive fork*, which satisfies the PCC, and the *interactive fork*, which does not. Their difference, however, is invisible in event-based causal modeling, which represents both scenarios with the same triad, $A \leftarrow C \rightarrow B$. What distinguishes the two types of fork, according to Salmon, is the process that connects these events (Salmon, 1984, p. 169). The conjunctive fork consists of separate and distinct processes that share specific background conditions (the fraternity file and the contaminated food). Conversely, the correlation in the interactive fork results from the spatiotemporal intersections of processes (the cue ball and the 8-ball). Based

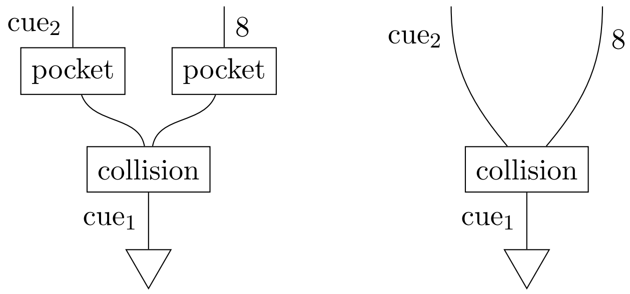

Figure 5: String diagrams for interactive forks. The diagram on the right is a simplified version of the left model, in which whether or not the balls fall is fully determined by their motion after the collision. cue_1 and cue_2 respectively represent the cue ball before and after the collision.

on this observation, Salmon concludes that determining whether the PCC holds for a given common cause structure requires attention to the underlying processes.

The categorical framework allows us to explicate Salmon's intuition and formulate two types of forks as distinct string diagrams. We begin with the interactive fork. An interactive fork is either a y-type splitting of one process into two or an x-type interaction where one of the inputs is abstracted away. For example, the billiard case consists of the interaction between the cue ball and the 8-ball, but only the cue ball is singled out as the cause of the collision. The left string diagram of Fig. 5 models this situation. In this model, the cue ball is set in motion by the state (represented by the downward triangle) and then collides with the 8-ball. This collision, depicted as the large box in the middle, imparts motion to both the cue ball (left wire) and the 8-ball (right wire). Depending on their momenta, they end up in pockets; this falling can be considered interactions between the balls and the pockets, denoted by the small boxes. Alternatively, if we consider that pocketing is determined entirely by each ball's motion after the collision, we can omit the small boxes and depict the situation as shown in the right diagram of Fig. 5. We use this simplified model below.

To see whether this structure leads to a violation of the PCC, let us now consider a functor that gives a probabilistic interpretation of the diagram. We assume that all processes are binary, with $F(cue_1) = F(cue_2) = F(8) = \{+, -\}$ representing whether the balls are hit or not hit (for cue_1) or pocketed or not (for cue_2 and the 8 balls). Then $F(collision)$ is given by the following 4×2 stochastic matrix:

	cue_1^+	cue_1^-
$cue_2^+, 8^+$	a	e
$cue_2^+, 8^-$	b	f
$cue_2^-, 8^+$	c	g
$cue_2^-, 8^-$	d	h

where each column gives the conditional probabilities of the outcomes when the cue ball is hit or not hit, respectively, so that $a + b + c + d = e + f + g + h = 1$.

Now, as Cartwright (1999, p. 8) observes, the PCC holds if and only if $ad = bc$ and $eh = fg$. But the functorial assignment imposes no such constraint: as long as each column sums to one, the outcome probabilities can take any value from 0 to 1. For instance, in the billiard example, one may assign $a = d = 0.5$ and $b = c = 0$, while $e = f = g = 0$ and $h = 1$. That is, when one hits the cue ball, both balls either fall into the pockets or remain on the table, each with a probability of 0.5; but when the cue ball is not struck, they both remain on the table, as expected. With this assignment, hitting the cue ball or not does not screen-off one outcome from the other. In general, multiple-output boxes such as the one in Fig. 5 admit any stochastic matrix, including those that do not satisfy the above screen-off condition. In this respect, they capture the intuition underlying the interactive fork—that interactions of two processes may result in a distribution that violates the PCC (Salmon, 1984; Cartwright, 1999).

Next, let us turn to the conjunctive fork. According to Salmon, a conjunctive fork consists of two non-overlapping processes, say A and B, that are independently affected by the same condition, C. Thus, the common cause C 'acts' or 'is used' twice, so to speak: once influencing A and once influencing B. In the plagiarism example, the same fraternity file was used independently by two students, each without affecting the other. In the case of food poisoning, the contaminated food was distributed to two or more guests, each of whom independently consumed their portion and became ill. In a string diagram, such situations are modeled using a special unit called the *copier*, which is denoted by a circle and serves to duplicate its input. Using this copier, the conjunctive fork can be modeled as in Fig. 6. In this model, the copier cp_C, represented by the circle in the middle, duplicates process C into two outputs, which then pass through distinct boxes f and g to yield A and B, respectively. In the plagiarism example, C represents the fraternity file, the copier symbolizes literal xeroxing or file copying, and the small boxes represent submissions by students. In the food poisoning incident, the copier would symbolize the serving from the contaminated pot, while the boxes represent the digestive processes of the guests.

It is important to note that a copier does not necessarily represent a physical duplication of the given process. Serving food from a pot onto plates does not constitute physically copying the food; rather, it prepares multiple processes (dishes) that operate independently with equal causal efficacy (of poisoning guests). In this sense, a copier represents the duplication of a process's "causal oomph" rather than of the process itself, indicating that its multiple causal actions operate independently of each other. The appropriateness of using a copier thus depends on the nature of the causal influence being modeled. The use of a copier to model simultaneous food poisoning rests on the assumption that serving plates from the same pot does not diminish the poisonous effect of either the individual plates or the pot itself. If, instead, our concern were with calorie intake, such an assumption would no longer be appropriate. Hence, although a copier does not presuppose the existence of any special mechanism that physically duplicates objects, it does represent a certain fact about the causal process under consideration—namely, that it operates independently on each of its causings.

To evaluate a distribution that arises from this model, we need to interpret a copier, that is, to specify the corresponding stochastic matrix $F(cp_C)$. Duplicating the input process means that, for an input c, a copier returns the pair (c, c). In probabilistic terms, this implies $P(c, c|c) = 1$, and zero otherwise. Hence, the stochastic matrix of a copier (of a process with n values) is the $n^2 \times n$ matrix $\Delta = (\delta_{(i,j)k})_{(i,j)k}$, where $\delta_{(i,j)k} = 1$ if $i = j = k$ and zero otherwise.

Without loss of generality, we again assume that all processes are binary with $F(A) = \{a_1, a_2\}, F(B) = \{b_1, b_2\}$ and $F(C) = \{c_1, c_2\}$. The stochastic matrices for boxes f and g are then 2×2 matrices, $F(f) = (f_{ik})_{ik}$ and $F(g) = (g_{jk})_{jk}$ with $i, j, k = \{1, 2\}$. With this setup, the conditional probability distribution $P(A, B|C)$

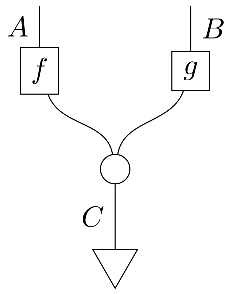

Figure 6: String diagram of a conjunctive fork.

is calculated by a matrix multiplication:

$$\begin{pmatrix} p(a_1,b_1|c_1) & p(a_1,b_1|c_2) \\ p(a_1,b_2|c_1) & p(a_1,b_2|c_2) \\ p(a_2,b_1|c_1) & p(a_2,b_1|c_2) \\ p(a_2,b_2|c_1) & p(a_2,b_2|c_2) \end{pmatrix} = F(f) \otimes F(g) \cdot F(cp_C)$$
$$= \begin{pmatrix} f_{11}g_{11} & f_{11}g_{12} & f_{12}g_{11} & f_{12}g_{12} \\ f_{11}g_{21} & f_{11}g_{22} & f_{12}g_{21} & f_{12}g_{22} \\ f_{21}g_{11} & f_{21}g_{12} & f_{22}g_{11} & f_{22}g_{12} \\ f_{21}g_{21} & f_{21}g_{22} & f_{22}g_{21} & f_{22}g_{22} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} f_{11}g_{11} & f_{12}g_{12} \\ f_{11}g_{21} & f_{12}g_{22} \\ f_{21}g_{11} & f_{22}g_{12} \\ f_{21}g_{21} & f_{22}g_{22} \end{pmatrix}. \quad (3)$$

Eqn. (3) satisfies the screen-off condition, because at each column $k = \{1, 2\}$, we have $f_{1k}g_{1k}f_{2k}g_{2k} = f_{1k}g_{2k}f_{2k}g_{1k}$. Indeed, noting that $f_{ik} = p(a_i|c_k)$ and $g_{ik} = p(b_i|c_k)$, Eqn. (3) shows that $p(a_i, b_j|c_k) = f_{ik}g_{jk} = p(a_i|c_k)p(b_j|c_k)$, namely that A and B are conditionally independent given C. Hence whenever the branching is made via a copier, the cause process screens off branches from each other.

Taken together, the above arguments support Salmon's claim that the two types of common cause structures—conjunctive and interactive—are characterized by underlying processes and interactions, and that only the former guarantees the satisfaction of the PCC. The validity of the common cause principle—as well as its generalized version, the causal Markov condition—has been a central point of contention in discussions about causal Bayesian networks (e.g., Cartwright, 1999; Hausman and Woodward, 1999, 2004; Sober, 2001; Näger, 2021). In response to counterexamples, proponents of the PCC have offered two types of rebuttals. One strategy is to question the validity of examples, claiming that they are artifacts arising from a misspecification of the model, such as a failure to incorporate the real common cause (Spirtes et al., 1993; Hausman and Woodward, 1999). However, without an independent proof of the PCC, this line of argument amounts to a devil's proof, urging the opponents to demonstrate the absence of the screening-off common cause. The second type of response acknowledges the existence of counterexamples and introduces additional notation to the DAG formalism that specifically denotes non-screening-off common causes (Schurz, 2017; Gebharter and Retzlaff, 2020). However, such responses remain ad hoc adjustments and fail to explain why common causes come in two different forms. In contrast, the process-based formalism not only accommodates non-screening-off causes but also offers a formal explanation for why the PCC holds in a specific type of fork structure—namely, those involving the copying or duplication of the common cause.

The discussion in this section naturally raises a question regarding the comparative representational capability between the category-theoretic framework and

Bayesian networks. It is known that Bayesian networks on DAGs correspond bijectively to categorical models in which common causes are represented using copier morphisms (Jacobs et al., 2019). This suggests that full-fledged categorical causal models, including those involving interactive forks, have strictly greater representational power than Bayesian networks. A detailed analysis of this point, however, must be deferred to future investigations.

7 Processes, Interactions, and Events

The preceding sections of this paper focused on the explanatory purport of process theory, analyzing it through category-theoretic modeling of etiological and constitutive explanations (Section 3), probabilistic interpretations (Section 4), explanatory relevance (Section 5), and the distinction between conjunctive and interactive forks (Section 6). This section shifts the focus to its metaphysical agenda and scrutinizes its ontological picture that "take[s] processes rather than events as basic entities (Salmon, 1984, p. 139)" from a category-theoretic standpoint. We do so by formulating and refining the basic ontological notions of the process ontology—namely processes, interactions, and events—within the categorical framework. This will substantiate the metaphysical intuition of process theorists that causal relationships are constituted by processes and their interactions, and also allow us to address significant challenges, including the problem of the identity of a process.

Let us first reconsider the categorical representation of processes. In Section 3, we identified causal processes with wires, which are usually interpreted as identity morphisms. For any object A, its identity morphism $1_A : A \to A$ is a morphism that does not change anything about A. More precisely, for any morphism $f : A \to B$, identities 1_A of A and 1_B of B satisfy $f = f \circ 1_A = 1_B \circ f$. Representing processes with identities is prima facie plausible because it satisfies the two major criteria of processes, namely the one based on mark transmission (Salmon, 1984) and the other on conserved quantity (Dowe, 2000). According to Salmon (1984, p.148), a genuine causal process is capable of transmitting a mark introduced by an external intervention, which, in the categorical framework, can be represented as a diagram surgery that replaces an unmarked state s_A of object A to a marked state s'_A (Section 5). Then identity morphism $1_A : A \to A$ trivially satisfies the mark transmission criterion, because $F(1_A) \circ F(s'_A) = F(s'_A)$ for any functor F, or in other words, marked state $F(s'_A)$ is transmitted intact through the process in the absence of additional intervention or interaction. Essentially the same argument shows that identity morphisms also accommodate Dowe's idea, according to which a causal process is "a world line of an object that possesses a conserved quantity (Dowe, 2000, p. 90)," like mass-energy, linear momentum, or charge. Let x denote a distribution over such quantities, e.g., the linear momenta of a ball thrown by a pitching machine.[5] We then again have $F(1_A) \circ x = x$, that is, the quantities are conserved throughout the process.

However, equating processes with identity morphisms turns out to be too restrictive for the purpose of modeling the dynamic aspects of causal processes. This is because, while identity morphisms by definition preclude any possibility of change, most actual processes that bear causal impacts are not strictly unchanging in every aspect. A moving ball incessantly changes its momentum and surface temperature

[5]We assume that our pitching machine has only a finite number of options for pitching speeds. To handle an infinite range of momenta, one must consider a functor to a broader category of standard Borel spaces and measurable maps (Fritz, 2020, sec. 4).

due to friction with the ambient air, no matter how slight these changes may be (Salmon, 1994, p. 309). This is even more true for organisms and other organic matter, as their identities are maintained only by the constant functioning of homeostatic mechanisms (Dupré, 2021). Identity morphisms, however, fail to capture the intuition that a single, identical process or object may change in certain aspects. Indeed, the dynamic nature of a process has been a source of metaphysical difficulty for process theorists themselves. On what ground can we say that one process remains the *same*, despite all the actual and possible changes it undergoes? This problem of the identity of a process has remained an open question in the philosophical discussion of process causality (c.f. Dowe 2000, p. 107; Salmon 1997, pp. 466-9).

The problem resolves once we note that the identity of a causal process emerges as an abstraction that disregards differences among its temporal or spatial parts. That is, a given spatiotemporally extended process is considered the same causal process despite changes occurring along its course, as long as these changes are negligible for modeling purposes. To implement this idea, we define the notion of *quasi-identities*: a morphism $f : A \to B$ of category C is a *quasi-identity with respect to functor* $F : \mathsf{C} \to \mathsf{D}$ if $F(f) = 1_{F(A)}$. Thus defined, quasi-identities include identities, as $F(1_A) = 1_{F(A)}$, but they also encompass a broader range of non-identity morphisms that represent substantive changes in domain category C, provided these changes are evaluated as negligible and reduced to identities in codomain category D. In the context of this paper, where D is $\mathsf{FinStoch}$ and $1_{F(A)}$ is the identity matrix of rank $|F(A)|$, a morphism $f : A \to B$ is considered a quasi-identity if it does not alter the probability distribution over the values $F(A)$ of A under consideration. For example, suppose that the trajectory of a ball from one time point to another is expressed by a morphism $f : A \to B$, but we are only interested in whether the ball is moving or stationary, such that $F(A) = F(B) = \{\text{moving}, \text{stationary}\}$. While the morphism may induce some changes, such as a slight reduction in momentum due to air friction, if the interval is short enough, these changes would not result in a qualitative difference in the ball's motion. Thus, f can reasonably be regarded as a quasi-identity representing one and the same continuous process. In this way, quasi-identities capture the Janus-like feature of a causal process by the functorial relationship between two categories.

We thus let quasi-identities represent causal processes and, accordingly, interpret wires in a string diagram as quasi-identities.[6] The proposed modification brings to the fore the role of idealization in causal modeling. Any modeling involves idealization, which regards inherently different things or phenomena as the "same." An appropriate idealization depends on the purpose of causal modeling (Salmon, 1994, p.309). The speed of a ball thrown by a pitcher may well be considered constant in a typical game situation; but when designing and improving baseball equipment, considering how friction and air resistance affect the ball's movement can be crucial. Therefore, it does not make much sense to ask whether something is a process or an interaction *tout court*—it all depends on the modeling purpose. The notion of quasi-identity accommodates such a model-dependent nature of causal processes by making explicit the role of modeling functor in their formulation.

Now let us move on to the second item of the ontological inventory of process theorists: interactions. As noted in Section 3, interactions are represented by boxes

[6] Our proposal can be seen as a generalization that bases string diagrams on pairs of monoidal categories and monoidal functors $\langle \mathsf{C}, F : \mathsf{C} \to \mathsf{FinStoch} \rangle$. The conventional interpretation corresponds to the special case where the monoidal functor is the identity functor $\langle \mathsf{C}, Id_{\mathsf{C}} : \mathsf{C} \to \mathsf{C} \rangle$, in which case quasi-identities coincide exactly with identities.

or (non-identity) morphisms. For instance, an x-type interaction between two causal processes is depicted by a box with two input and two output wires. This, however, is not sufficient to capture Salmon's CI (causal interaction) criterion (Salmon, 1984, p.171), which requires x-type causal interactions to yield mutual modifications of the characteristics of the processes involved. Not all multi-pronged boxes imply such an interplay. For instance, a parallel product of two I-type boxes, $g : A \to B$ and $h : C \to D$, gives rise to a box with two inputs and two outputs, $g \otimes h : A \otimes C \to B \otimes D$ (see (1) in Section 3); however, its two inputs undergo independent processing without affecting each other.

A straightforward way to avoid this issue is to define an (x-type) interaction as box/morphism $f : A \otimes C \to B \otimes D$ that is not decomposable into a parallel products of two boxes/morphisms $g : A \to B$ and $h : C \to D$ such that $f = g \otimes h$. However, here again we must take into account the fact that what counts as interactions depends on the modeling context. Do two intersecting sound waves coming from different sources interact with each other? In everyday conversation settings, we can generally treat them as mutually independent. But a sound engineer working for a concert hall may well take it as an interaction and be concerned with its acoustic effect. Such model dependency can again be accounted for by defining interactions in terms of a functor. Specifically, we call a morphism $f : A \otimes C \to B \otimes D$ of C an x-type interaction with respect to a functor $F : \mathsf{C} \to \mathsf{FinStoch}$ if $F(f)$ is not decomposable into two stochastic matrices, namely there are no two stochastic matrices, M_1 of size $|F(B)| \times |F(A)|$ and M_2 of size $|F(D)| \times |F(C)|$, such that $F(f) = M_1 \otimes M_2$. This condition imposes an interplay between the two processes, in the sense that there is an intervention on at least one of them that makes a difference in the counterpart.[7]

Interactions represent changes within the processes involved. Since changes occur over a certain interval, we believe that most, if not all, interactions extend through time and space. Accordingly, neither processes nor interactions are instantaneous in our framework. This conclusion follows naturally from the compositional nature of string diagrams. As discussed in Section 3, sequentially stacked boxes can be viewed as a single, larger box. Through repeated composition, any string diagram can ultimately be interpreted as a single unified box. If the interaction represented by a box were instantaneous, it would imply that all causal changes complete in an instant—a clearly absurd conclusion. Moreover, because interactions are spatiotemporally extended, they can aptly represent gradual changes, such as deceleration of a ball's velocity due to air resistance. A key advantage of process theory lies in its ability to directly represent such continuous changes, without reducing them to a finite or infinite chain of instantaneous events—thereby avoiding Zenoesque paradoxes, as we will discuss shortly.

We have thus located the two main ingredients of the process ontology, causal processes and interactions, within the categorical framework. Finally, let us see how events fit in this picture. Compared to processes, "events are relatively localized" and "represented by points" rather than lines in the space-time manifold (Salmon, 1984, p. 139). But to what extent should events be localized? Although Salmon lists a sneeze, a baseball colliding with a window, and the activation of a photocell by a pulse of light as examples of events, all such actions extend both in space and

[7]Salmon's CI condition (Salmon, 1984, p. 171) requires that an intervention on either of the two inputs of an x-type interaction brings about a difference in the other. However, this requirement is unduly strong, precluding cases where one of the processes involved remains invariant during the interaction, as in the case of a catalyst in a chemical reaction. Our proposal here weakens Salmon's original condition to include such cases as x-type interactions.

time and cannot be completed "at a point." A genuine event can only be found in the limit where the time interval approaches zero. At such an extremity, any change would vanish, with all characteristics of the objects involved (the person who sneezes, the ball, the photocell, etc.) being self-identically determined.

Events in this strict sense are best modeled with identity morphisms. In our parlance, identities are a trivial case of quasi-identities, which are mapped to identities by any functor. From a philosophical perspective, this amounts to regarding events as "degenerated" processes that preclude any moment of change. Specifically, for an object B we let its identity 1_B represent an event that captures the state of B at any given moment. Pre-composed with another morphism $f : A \to B$ as $1_B \circ f$, the event 1_B represents the state of B affected by A through f. Likewise, 1_A in $f \circ 1_A$ indicates the state of A that causally affects B. Taken together, a morphism sandwiched between identities, as in $f = 1_B \circ f \circ 1_A$, can be read as event 1_A affecting event 1_B via f. In this sense, events or identities serve as causal relata.

Two caveats are in order, however. First, events or identities are neither the only nor the typical causal relata in our framework. Causal relata of a given causal relationship f are the morphisms that are pre- or post-composed with f. Since these morphisms can be non-identities, processes (i.e., quasi-identities) and even interactions can also serve as causal relata. Second, causal relata do not determine the causal relationship. There may be another morphism $f' : A \to B$ with $f' \neq f$, which, when similarly composed as $1_B \circ f' \circ 1_A$, represents an alternative way in which 1_A influences 1_B. This is contrasted with the DAG formalism, where the distinction among causal pathways running from the same cause to the same effect can be made only by way of intermediate nodes, such as $A \to C \to B$ and $A \to D \to B$. This is because the identity of a causal edge, defined as a pair of nodes, is determined solely by its relata. In contrast, causal relationships in process theory stand on their own, while events are ancillary attendants arising only as cross-sections of a causal process.

This minimalist conception of events explains why process theory is free from the Zenoesque paradox, which asserts the impossibility of constructing substantive changes or the propagation of causal influences from instantaneous events (Russell, 1912; Hitchcock, 2004). In our framework, causal production and propagation are represented by morphism $f : A \to B$ and quasi-identity $g : A \to A$, respectively. Consider causal production $f : A \to B$. By introducing an intermediate event 1_C, f can be split into two successive parts such that $f = f_2 \circ f_1 : A \xrightarrow{f_1} C \xrightarrow{f_2} B$. Further decompositions would result in $f = f_n \circ \cdots \circ f_1$ with $n-1$ intermediate events $1_{C_1}, \ldots, 1_{C_{n-1}}$. However, the entire causal relationship $f : A \to B$ cannot be reduced to these events, for if $A \neq B$, the decomposition $f_n \circ \cdots \circ f_1$ must include at least one non-trivial morphism to constitute a morphism from A to B. Even in the case of a quasi-identity $g : A \to A$ with $g \neq 1_A$, the decomposition must contain at least one non-trivial morphism, as identities (of A) alone sum only to an identity. Thus, in either case, there can be no substantive causal chain consisting solely of events. From the categorical perspective, the Zenoesque paradox simply reflects a truism: a non-identity morphism cannot be reduced to a collection of identities.

This last point highlights the fundamental difference between event-based and process-based ontologies. In the former, the basic building blocks of the world are static events, while anything that involves dynamic changes, including causal relationships, has only a derivative status and must be constructed out of events. In contrast, process-based theory takes processes to be primitive and events to be derivative. Indeed, events are conceptualized as "degenerated" morphisms (i.e., identities) that may appear only at the endpoints of another morphism. They are like cross-sections with zero thickness of a candy stick. Cross-sections come into existence only by cutting the candy, and one can never make the whole candy by gluing its cross-sections together. The ontological priority of processes over events can be understood in a similar way, and this intuition of the process ontology is best captured by morphisms.

8 Conclusion

Ever since Hume, most philosophical discussions of causality have been framed within the event-based ontology, analyzing causation as a special kind of relationship between events. In opposition to this standard view, process theorists have sought to depict causality in terms of processes and their interactions, but they have lacked a mathematical framework to formally implement their ideas.

This paper introduced one such formalism based on the recent development of categorical causal modeling (Jacobs et al., 2019; Fritz and Klingler, 2022; Lorenz and Tull, 2023). The categorical framework formulates the basic building blocks of the process ontology—namely processes, interactions, and events—as components of string diagrams, as well as its two major explanatory strategies—etiological and constitutive explanations—by way of the compositional features of monoidal categories. Moreover, combined with a functor, it enables the probabilistic and counterfactual analysis of causality, thereby addressing the problem of explanatory irrelevance and providing a means to express the difference between conjunctive and interactive forks.

The new formalism rekindles philosophical debates over the nature of causality. Process theory adopts a richer ontology than event-based approaches, recognizing not only events but also processes as fundamental constituents of reality. For Salmon (1984), this ontological enrichment was driven by its explanatory purchase: he argued that it is precisely by focusing on underlying processes that we can account for the crucial distinction between conjunctive and interactive forks. The category-theoretic modeling gives substance to this ontological intuition by providing a formal and unified perspective on processes and interactions, representing them as morphisms. Events, in contrast, are treated as degenerate processes, namely identity morphisms. Our discussion in Section 6 supports Salmon's claim by distinguishing conjunctive and interactive forks through two types of morphisms—those that involve a copier and those that do not. The bijective correspondence established by Jacobs et al. (2019) between Bayesian networks and categorical causal models that involve only conjunctive common causes—namely, those in which all y-type interactions are mediated by a copier—corroborates this point, demonstrating that models with interactive forks lack counterparts in the framework of Bayesian networks. Salmon, therefore, was right: the process ontology indeed offers an explanatory advantage.

As in this case, the Principle of Common Cause has been an epicenter in philosophical debates about causation (e.g. Cartwright, 1999; Hausman and Woodward, 1999, 2004; Sober, 2001). Although some proponents of the principle have tried to derive it as an intrinsic feature of the notion of causality (Hausman and Woodward, 1999), our discussion in Section 6 makes it clear that it is actually a postulate or limitation of the DAG formalism, while the categorical framework presumes no such restriction on the common cause structure. A natural question that follows is when the PCC holds and when it fails—that is, what conditions license the use of a copier. In quantum mechanics, it is known that the use of a copier is not appropriate in situations involving arbitrary unknown quantum states, due to the no-cloning theorem. Can we likewise hope for conditions that distinguish between conjunctive and interactive common causes in macroscopic, non-quantum settings? In this way, the categorical framework provides a fresh perspective on longstanding philosophical debates surrounding the PCC.

The present paper has focused on the philosophical merits of process theory, particularly its capacity to offer a more encompassing view of causation. An equally important question concerns its methodological advantages: does it allow for more flexible or effective modeling of diverse causal structures? Recent work suggests that the categorical framework is well-suited for formally defining abstraction between causal models of different granularities, via natural transformations between causal models (i.e., functors) (Rischel and Weichwald, 2021; Otsuka and Saigo, 2022). On the other hand, the DAG formalism enjoys a wealth of methods for estimating causal structures from observational data (Spirtes et al., 1993; Peters et al., 2017). To what extent these results are transferable to each other is a task for future investigations. In conclusion, we hope that the mathematical formulations of process causality discussed in this paper will encourage further theoretical work aimed at integrating these two perspectives, with the goal of combining their strengths and advancing a more comprehensive understanding of causality.

Acknowledgment

We are grateful to Tobias Fritz and two anonymous referees for their valuable feedback, which significantly helped improve the manuscript.

References

Bechtel, W. (2006). *Discovering cell mechanisms: The creation of modern cell biology.* Cambridge University Press.

Bechtel, W. and Abrahamsen, A. (2005). Explanation: a mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36(2):421–441.

Cartwright, N. (1999). *The dappled world.* Cambridge University Press.

Coecke, B. and Kissinger, A. (2017). *Picturing quantum processes.* Cambridge University Press.

Craver, C. F. (2001). Role functions, mechanisms, and hierarchy. *Philosophy of Science*, 68:53–74.

Dowe, P. (2000). *Physical causation.* Cambridge University Press.

Dupré, J. (2021). Causally powerful processes. *Synthese*, 199(3):10667–10683.

Fong, B. (2013). Causal theories: A categorical perspective on Bayesian networks. arXiv:1301.6201 [math.PR].

Fritz, T. (2020). A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics. *Advances in Mathematics*, 370:107239.

Fritz, T. and Klingler, A. (2022) The d-separation criterion in categorical probability. arXiv:2207.05740 [math.ST].

Gebharter, A. and Retzlaff, N. (2020). A new proposal how to handle counterexamples to Markov causation à la Cartwright, or: fixing the chemical factory. *Synthese*, 197(4):1467–1486.

Glennan, S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, 44(1):49–71.

Glymour, C. (2004). Review of Making things happen: A theory of causal explanation, by J. Woodward. *The British Journal for the Philosophy of Science*, 55(4):779–790.

Hausman, D. M. and Woodward, J. (1999). Independence, invariance and the causal Markov condition. *British Journal for the Philosophy of Science*, 50:521–583.

Hausman, D. M. and Woodward, J. (2004). Modularity and the causal Markov condition: A restatement. *British Journal for the Philosophy of Science*, 55(1):147–161.

Hitchcock, C. (2004). Causal processes and interactions: What are they and what are they good for? *Philosophy of Science*, 71(5):932–941.

Jacobs, B., Kissinger, A., and Zanasi, F. (2019). Causal inference by string diagram surgery. In *Foundations of Software Science and Computation Structures*, pages 313–329. Springer International Publishing.

Lorenz, R. and Tull, S. (2023). Causal models in string diagrams. arXiv:2304.07638v1 [cs.LO].

Machamer, P., Darden, L., and Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1):1–25.

Näger, P. M. (2021). Evidence for interactive common causes. Resuming the Cartwright-Hausman-Woodward debate. *European Journal Philosophy of Science*, 12(1):2.

Otsuka, J. and Saigo, H. (2022). On the equivalence of causal models: A category-theoretic approach. In *Proceedings of the First Conference on Causal Learning and Reasoning*, pages 634–646.

Pearl, J. (2000). *Causality: Models, reasoning, and inference.* Cambridge University Press.

Peters, J., Janzing, D., and Schölkopf, B. (2017). *Elements of causal inference: Foundations and learning algorithms.* The MIT Press.

Rischel, E. F. and Weichwald, S. (2021). Compositional abstraction error and a category of causal models. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, PMLR 161:1013-1023.

Russell, B. (1912). On the notion of cause. *Proceedings of the Aristotelian Society*, 13:1–26.

Salmon, W. C. (1984). *Scientific explanation and the causal structure of the world.* Princeton University Press.

Salmon, W. C. (1994). Causality without counterfactuals. *Philosophy of Science*, 61(2):297–312.

Salmon, W. C. (1997). Causality and explanation: A reply to two critiques. *Philosophy of Science*, 64(3):461–477.

Schurz, G. (2017). Interactive causes: Revising the Markov condition. *Philosophy of Science*, 84(3):456–479.

Sober, E. (2001). Venetian sea levels, British bread prices, and the principle of the common cause. *British Journal for the Philosophy of Science*, 52(2):331–346.

Spirtes, P., Glymour, C., and Scheines, R. (1993). *Causation, prediction, and search.* The MIT Press.

Woodward, J. (2003). *Making things happen.* Oxford University Press.