Foundations of Physics (2025) 55:63
https://doi.org/10.1007/s10701-025-00875-w

RESEARCH

Causal Set Theory is (Strongly) Causal

Sam Baron[1] · **Baptiste Le Bihan**[2]

Received: 1 August 2024 / Accepted: 30 June 2025

Abstract
Causal Set Theory (CST) is a promising approach to fundamental physics that seems to treat causation as a basic posit. But in exactly *what* sense is CST causal? We argue that if the growth dynamics is interpreted as a physical process, then CST employs relations of actual causation between causal set elements, whereby elements bring one another into existence. This is important, as it provides a better sense of how CST works, highlights important differences from general relativity—where relations between spacetime points are typically seen as cases of mere causal connectibility rather than actual causation of the relevant type—and points toward a specific understanding of the emergence of spacetime within CST.

Keywords Causation · Causal · Causal set theory · Spacetime emergence · Quantum gravity

1 Introduction

Causal Set Theory (CST) is a promising approach to developing a quantum theory of gravity.[1] It aims to reconstruct physics from a discrete structure of elements and relations, which collectively give rise to the macroscopic world as we know it, including the spatiotemporal manifold described by general relativity. The basic structure of CST is a causal set: a discrete set of elements C connected by an ordering relation commonly denoted as $\prec$. The structure is usually described as evolving through the operation of a dynamical law that adds elements one-by-one.

[1] [4, 5, 7, 8, 19, 25, 26].

✉ Sam Baron
s.baron@unimelb.edu.au

[1] University of Melbourne, Melbourne, Australia

[2] University of Geneva, Geneva, Switzerland

Published online: 12 July 2025

The causal aspect of CST is often emphasised by the engineers of the theory (and also by philosophers, see e.g. [16, 40]). For instance, Dowker writes that

> CST arises by combining discreteness and causality to create a substance that can be the basis of a theory of quantum gravity. [6, p. 446]

Rideout and Sorkin describe CST in similar, causal terms:

> ... the order relation constituting C [a causal set] will be causal in the dynamical sense, and not only in name. [26, p. 6]

Sverdlov and Bombelli also view causal set theory as causal, writing:

> The causal set approach to quantum gravity uses causal sets as the only fundamental structure for the description of the geometry of spacetime. [30, p. 1]

But in exactly *what* sense is CST causal? Our goal is to answer this question. Focusing on the standard sequential growth dynamics for CST, we argue that if the dynamics is treated as a real physical process, then CST makes use of causation of a specific sort, whereby causal set elements causally depend for their existence on previous elements. To show this, we use the framework of interventionism [22, 33], which is a helpful tool for identifying causal relations of various types. Treating the growth process as a real physical process is controversial (see [1, 15, 38]). As we shall see, however, our analysis might help to reimagine the growth process by treating it as a causal process rather than a process of temporal passage or becoming. This is significant, since much of the resistance to treating the growth process as a real process stems from the conceptualisation of it as a process of becoming.

We begin, in §2, by clarifying and motivating the project of examining causation in CST. In §3 we outline CST in a bit more detail and in §4 we summarise the interventionist framework developed by [22] and [33]. After that, in §5, we apply interventionism to the dynamics of CST and show that causal set elements causally depend for their existence on previous elements. In §5.1 we assume that the classical sequential growth dynamics corresponds to a real physical process. In §5.2 we modify this assumption in line with [38] and show that even taking this modification into account, interventionism identifies causal dependence between elements of the relevant sort. §5.3 looks at objections and §6 sums up.

2 Clarification and Motivation

We begin by clarifying the core question at hand: in what sense is causal set theory causal? This question stands in need of clarification because there is already a well-understood sense in which causal set theory is causal.

To see this, it is instructive to consider general relativity (GR) and the way the concept of causality is deployed therein. In the framework of general relativity, causation usually refers to the existence of a lightcone structure that places constraints on what

can, and cannot, be dynamically connected via propagating signals. Two events in a manifold are causally related in this sense if and only if there is a timelike or null curve connecting them. It is easy to understand why one would want to qualify such a structure as being causal, as it specifies the domain of possible, physical *interactions*. However, this structure only delimits the domain of entities that can or cannot be connected by physical signals, which could be consistent, in principle, with the view that no stronger notion of causation is required by general relativity.

This notion of causation is important to CST and is the one from which the research program inherited its name. For CST was founded on a set of theorems developed by [12, 20] and others, showing that the metric structure of spacetime can be derived from the lightcone structure of general relativity, up to a conformal factor. Now, it could be that the only sense in which CST is causal is that it involves causation in the sense of recovering the causal structure of general relativity. The point was nicely made by Wüthrich and Huggett who write:

> The worry is familiar from relativity, and surely from the causal theory of (space)time, to which pedants (ourselves included) have always interjected that the 'causal structure' of spacetimes merely captures a minimally necessary, but not sufficient, connection between events for them to be causally related as cause and effect. The objection is motivated by the observation that we do not attribute causal efficacy to all timelike or null relations; given an event, we take neither all events in its past lightcone to be its 'causes', nor all events in its future lightcone to be its 'effects'. [40, p. 14]

We grant that CST is causal in the sense described above. What we want to know is whether it is causal in a stronger sense. We can sharpen the point, initially, by differentiating between two notions of causation. The first is the notion identified above, of causal connectibility. We consider this to be a *weak* notion of causation. That's because, to say that there is a relation of causal connectibility between x and y is just to say that there is a possibility of causal influence between x and y. This says nothing of any actual causal relations.

A stronger notion of causation goes beyond causal connectibility to actual causation. Thus, x and y are causally connected in this stronger sense when there is an actual causal relation between them. That is, when x actually does—not merely could—influence y.

Note that causal connectibility does not imply the presence of actual causation. We will demonstrate this further later on. For now, the point can be illustrated by considering a general relativistic spacetime that is entirely empty—no matter fields (formally, all matter fields one could introduce are zero-valued everywhere), only the metric field is present. The metric allows one to define causal connectibility (via timelike or null curves). Yet, in such a scenario, no actual causal processes are occurring, since there are no entities undergoing change or exerting influence via the matter fields. This shows that causal connectivity can, in principle, exist absent actual causation.

Thus, when we ask whether causal set theory is causal, what we want to know is whether it is causal in the stronger sense of actual causation, which goes beyond causal connectibility. However, this is not yet enough to specify the focus of this

paper. Our interest is not just in *any* type of actual causation. Rather, we are interested in actual causation of a very specific variety. For, as we see it, there are two distinct types of actual causation.

Think again of the general relativistic case. One type of actual causation might hold between whatever is located at two spacetime points. So, for instance, we might look at how light emitted from point x affects a detector located in its future light-cone, at point y. Call this type of actual causation: *material causation*. Another, more radical type of actual causation might hold between the spacetime points themselves, regardless of what is located there. So, for instance, we might say that a spacetime point x actually causes spacetime point y to exist. Call this type of actual causation: *elemental causation*, since it holds between the basal elements of a theory, in this case between the spacetime points.[2]

Elemental causation is best demonstrated using the interventionist framework outlined below. The basic idea, though, is this: elemental causation occurs when, were one to 'remove' an element x this would make a difference to whether another element y exists, or to the probability of y existing. For instance, a spacetime point y depends on a spacetime point x in this sense when, hypothetically, *removing* x would require removing y as well: y's existence is not independent of x's.

Now, while we have used general relativity to illustrate the notion of elemental causation, we are not claiming that there is causation of this type between spacetime points in general relativity (though we return to this in §5.3). Rather, the point is just to clarify the kind of causation we are interested in with regard to CST. In the case of CST, we are interested in elemental causation, and not material causation. What we want to know is whether the existence of a causal set element depends on the existence of elements in its past. Thus, our question is: *do causal set elements cause one another to exist and thus does CST exhibit elemental causation?*

This question matters for three reasons. First, answering it helps us to better understand how CST works. For it clarifies the role that causation plays under a specific interpretation of that theory, where the growth process is taken seriously (more on this in a moment), which in turn helps us to better understand what the world would have to be like for CST so interpreted to be true. At the fundamental level, there would need to be actual causal relations linking causal set elements. Where, and this is important, those actual causal relations hold between the elements themselves (rather than whatever is located at those elements), so that elements literally bring other elements into existence.

Second, it has been argued that CST is not, fundamentally, a spatial, temporal or spatiotemporal theory (see [36, 39]). Rather, spacetime emerges from a causal set. Moreover, this happens only under certain conditions. Many physically possible causal

[2] One might worry that elemental causation collapses into causal connectibility. Again, consider the case of general relativity. Both causal connectibility and elemental causation appear to exist merely in virtue of spacetime structure, without mediation by a matter field. Indeed, both connectibility and elemental causation could exist in a vacuum world. But does that entail the two notions are equivalent? We think not, as any coextension is at best contingent. For one can accept causal connectibility without endorsing the view that spacetime points cause other points to exist. Indeed, as we discuss in §5.3, this is precisely what we should think in the case of general relativity, which further highlights the way that the two notions of causation can come apart.

sets do not correspond to any solutions of general relativity and thus to any spacetime. It thus seems that the basic structure of a causal set—the structure that is present across all models of the theory—is not spatial, temporal or spatiotemporal structure.

Suppose that CST is indeed fundamentally non-spatiotemporal, with spacetime emerging (we will return to this assumption in a moment). If spacetime is emergent in CST it is clearly emergent from causation *in some sense*, since causation is one of the only basic posits of the theory. What we don't have a clear picture of just yet is *how* spacetime emerges from causation.

One possibility is that spacetime emerges in CST from relations of actual causation, analysed along interventionist lines, and consonant with elemental causation [2]. The viability of this picture of spacetime emergence thus hangs on the question of whether there is actual causation of a specific variety 'at the bottom', as it were in CST, that can be used as the basis for spacetime.

If there are no such actual relations of causation, then this tells us that a new metaphysical account of how spacetime emerges in CST is needed, one that relies on a different notion of causation, perhaps specified in terms of causal connectibility or a notion of actual causation other than elemental causation. Perhaps indeed such an account can be developed. The point, though, is that until we consider what type of causation is operative in CST, we have less available conceptual resources for producing a metaphysical account of how spacetime emerges.

A third reason for considering causation in the context of CST relates to broader questions of whether causation plays a role in physics. A number of philosophers maintain that causation plays no role in fundamental physics. Norton [21], for instance, regards the concept of causation as pertaining to an approximately true, but strictly false, folk theory, which has been falsified on many occasions with the development of science. Science, in some sense, transcends this notion of causation. As [27, p. 2] famously put it:

> The law of causality, I believe, like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no, harm.

In a similar vein, causal republicanism has gained momentum and is portrayed as the view that causation, although remaining a useful notion, is generated by agents and is not found in the fundamental ontology of physics [23].

The focus of this scepticism about causation is not on causal connectibility, which as indicated is present even in GR. The focus is on actual causation. The question is whether causation in any sense stronger than connectibility has a role to play in physics. If CST makes use of actual, elemental causation, then it might provide an interesting case study in this broader debate about the role of causation. If actual, elemental causation is no part of CST, then this would tend to strengthen the sceptical position that Norton, Price and Russell have expressed. It is thus worth looking a bit closer at CST to see what role causation plays in the theory.

With these three points in mind, we propose a limited answer to the question of whether there is elemental causation in CST. We argue for the affirmative, that there is elemental causation in CST, however our argument is conditional, in this sense:

there is elemental causation in CST on the condition that a particular interpretation of the dynamics is adopted. Specifically, an interpretation in which the growth process in the standard dynamics corresponds to a real physical process.

We grant that there is a controversy over whether the growth process should be taken as a real physical process at all. We consider this more fully in §5. For now, we note two things. First, some of the main engineers of CST (e.g., [6, 29]) take the growth process seriously, which is at least an invitation to do the same in philosophy, to see what can be made of it. Second, the controversy regarding the growth process revolves mainly around the claim made by some physicists and challenged by some philosophers that the process is one of *passage*, analogous to the passage of time found in theories like the growing block theory of time (see [1, 38]). We think it is exactly this feature of the controversy that makes our argument useful. For if there is causation between the elements of causal set theory, then this opens up an alternative way to interpret the growth process as physical. We can interpret it as the unfolding of a specific causal process, rather than a process involving the passage of time.

This has some potential advantages: thinking of the process as a causal process may provide a new way to take the process physically seriously. The passage of time is already controversial, and so interpreting CST in this way wraps it in this controversy from the beginning. Looking at the theory from a purely causal perspective may be an improvement. That's because the worries that philosophers have voiced concerning the growth process being physical seem less concerning with regard to causation. The main worry is that treating the process as a physical process results in a kind of metaphysical indeterminacy, which is unattractive if we are considering the process as temporal passage, which usually lacks this feature. However, as [14] argues, causation *in general* appears to be indeterminate, and so the kind of indeterminacy at issue may be exactly what we should expect, and thus no more of a problem than for causation in general. This is not to say that indeterminacy is not vexing. The point is that a causal interpretation of the growth process may move the bump in the carpet in a helpful way, by treating the problem of metaphysical indeterminacy as an issue for causation in general, rather than the metaphysical interpretation of CST *per se*.[3]

Our argument is useful for another reason: the specific type of actual causation we are focusing on gives us a clearer picture of the type of emergence we can expect (again, assuming the growth process is physical). As noted, the kind of causation we are interested in is one where the causal set elements actually cause one another to exist. Philosophically, this reveals the kind of metaphysical picture at issue: CST coupled to an interpretation of the growth process as physical is highly analogous to the causal theory of time advocated by [18] [1715] (1989), in which time is generated from a causal process. Scientifically, it reveals that the emergence of spacetime is the outcome of a causal process. This makes it analogous to the type of emergence we find in phase transitions, where a system undergoes a sequence of causal interactions

[3] Our approach can also help to provide a non-spatiotemporal interpretation of CST that is still compatible with growth. On this interpretation, growth is a process of causal, rather than spatiotemporal development. This is potentially useful as it may reconcile the conviction that some physicists have that the growth process is real, with the apparently non-spatiotemporal nature of the theory, something that cannot otherwise be done if growth is interpreted as passage.

that trigger a state-change. This provides an interesting lens through which to view causal set theory under a growth interpretation and, perhaps, more generally.

3 Causal Set Theory

To determine whether CST makes use of actual causation, we need a bit of set up. First, we need to provide an overview of CST and the standard sequential growth dynamics. Next, we need a framework for identifying actual causation to guide our search. For this, we use interventionism, since it is the leading approach to actual causation available. We outline interventionism in the next section. In this section, we outline some details of CST.

The kinematics of CST can be represented as a causal set, or *causet* for short:

$$\langle C, \prec \rangle$$

Here C is a set of elements and $\prec$ is a relation defined on C.[4] CST obeys three axioms. First, $\prec$ is a strict partial order over C: it is irreflexive, transitive and anti-symmetric. This turns each causal set into a partially ordered set (a poset).[5] Second, causal sets are locally finite:[6]

$$\forall a, c \in C, \mathrm{card}(\{b \in C | a \prec b \prec c\}) < \infty$$

Where card denotes the cardinality of a set, and $(\{b \in C | a \prec b \prec c\})$ denotes the set of elements that are between a and c by $\prec$. Local finiteness is a way of imposing discreteness on causal sets. The basic idea being that between any two causal set elements ordered by $\prec$ there is at most a countable, finite number of elements. Third, C is countable (that is, there exists an injective function from the set of elements to the natural numbers).

In addition to these axioms, we adopt further assumptions from Rideout and Sorkin [26]:

a. The *past* of an element $x \in C$ is the subset $\mathrm{past}(x) = \{y \in C | y \prec x\}$.
b. A *chain* is a linearly ordered subset of C.
c. An *antichain* is a totally unordered subset of C.
d. A *partial stem* of C is a finite subset which contains its own past.

[4] This relation is sometimes called the *relation of causal precedence*. However, as a referee notes, given that $\prec$ is transitive, it does not correspond to causation as many think of that notion. We don't see this as a problem for what follows. In line with the discussion in §1, the relation is supposed to be causal in the sense that it captures the light-cone structure of general relativity, i.e. the timelike relations, which are indeed closed under transitivity.

[5] In the most general statement of the theory, $\prec$ is reflexive. Here we follow Rideout and Sorkin ([26, p. 024002-2]) and impose an irreflexivity convention on $\prec$ to turn causal sets into strict partial orders. We do this because that is how they set up the growth dynamics.

[6] We use here a statement of local finiteness from [37].

e. A *link* of C (i.e., an instance of $\prec$) is irreducible (cannot be built from other relations on C).
f. A *path* in C is a chain of elements, each related by a link.

The dynamics of a causal set is defined in terms of a step-wise random, Markovian process of 'growth'. The process starts with the empty set. At each step of the process, a new element is added to the causal set with some probability. The new element is either connected to an existing element with a link, or is unconnected to existing elements.

The elements that are added through the growth process are given a labelling in terms of integers, such that $x \prec y \rightarrow \text{label}(x) < \text{label}(y)$. Importantly, no element is added into the past of an existing element. This condition is commonly called *internal temporality*. As Rideout and Sorkin note, the labels add an element of 'gauge' into the initial statement of the dynamics. The elements of causal set theory are treated as though they came into existence in a definite order (thus suggesting an external time for the growth process). Ultimately, though, conditions are added to the dynamics to render them gauge invariant (more on this in a moment).

We can represent the dynamics of causal set theory in terms of a sequence of causal sets that originates from the empty set and then terminates in a causal set *C*. This sequence can also be represented as a partially ordered set, though this ordering is not via $\prec$(which joins elements within a causal set) but by a second ordering, $\rightarrow$, which joins sets of elements. For causal sets *x* and *y* such that $x \rightarrow y$, we say that *x* is the *parent* and *y* is the *child*. The stochastic process can then be captured via a set of transition probabilities from parents to children: the probability of a given child, given a certain parent. The set of transition probabilities issuing from a parent causal set to its children should always sum to unity. Rideout and Sorkin [26] call this the *Markov Sum Rule*.

The transition probability a_n from a causal set C_n made of *n* elements to a causal set C_{n+1} is given via the following rule:

$$a_n = p^m(1-p)^{n-\varpi} \tag{1}$$

Where *p*, a fixed parameter, is the probability of a new element being added that is linked to an existing element; $q = (1-p)$ is the probability of a new element being added that is unlinked to existing elements; *m* is the number of maximal elements in the precursor set, which for the new element *e* added in C_{n+1} is the set of elements in C_n such that each member of the set is (i) ordered with respect to *e* by $\prec$; (ii) is prior to *e* in that ordering. The maximal elements of the precursor set are those elements in the set that are not prior to any element in the ordering $\prec$; ϖ is then the size of the entire precursor set.

To complete the dynamics, two further conditions are added. The first is *discrete general covariance*. The second is *Bell causality*. We take each in turn. Discrete general covariance is added to render the dynamics gauge invariant, thereby undoing the fact that the dynamics is generally written in terms of a gauge (via a labelling of the growth order of elements). The upshot of discrete general covariance, Rideout

and Sorkin note, is that the order of growth specified by the gauge is not physically significant. As they put it, "the labels carry no physical meaning" [26, 26

> If γ is any path through the poset $\mathcal{P}$ of finite causal sets that originates at the empty causet and terminates at C, then the product of the transitional probabilities along the links of γ must be the same as for any other path arriving at C.

To see the idea, let us consider a toy model. In this toy model, we are considering the growth of a causal set to its third element. We assume that the causal set continues to grow to infinity, but we don't model those stages of the process for obvious reasons. In the model, we also use a choice of label, and thus write the dynamics in a choice of gauge. For this, we use colours (causal sets with labels are often called 'coloured' causal sets in the literature). The colours represent the order of the growth elements chosen to produce the model. It is this information—the coloured labels—that discrete general covariance will ultimately factor out. But, as is standard, we start with a choice of labels first, and then apply discrete general covariance to produce the generally covariant dynamics. The toy model can be depicted as follows:

Notice that in Fig. 1 there are three paths to the following causal set:

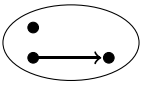

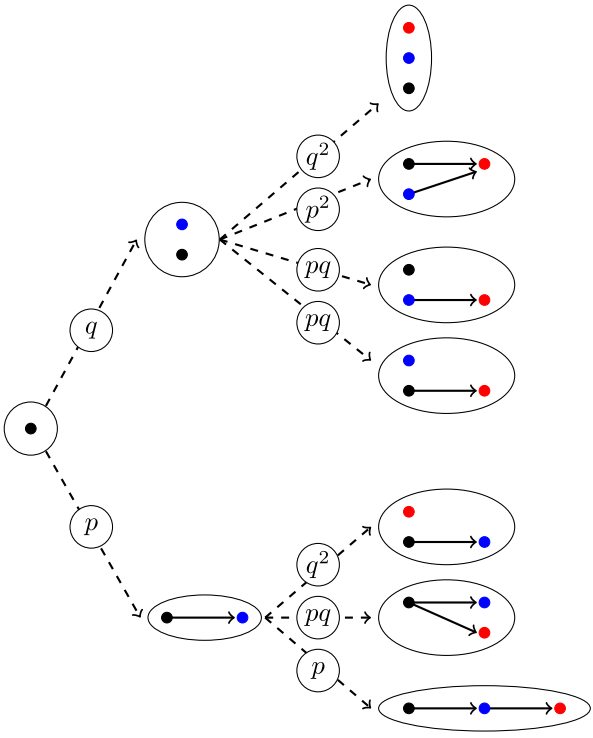

Fig. 1 A toy model of growth dynamics, where p is the probability of adding an element that is linked to existing elements; and $q = (1 - p)$ is the probability of adding an element that is unlinked to existing elements. Transition probabilities in the bubbles are derived from p and q using the general rule $a_n = p^m(1 - p)^{n-\varpi}$ introduced before

What discrete general covariance tells us, is that the product of the transition probabilities along the three paths that lead from the first causal set to this causal set should be the same. This is, in fact, the case in our diagram. There are three paths to the causal set in question. The probabilities can be computed as follows:

P1 $: q \times pq = pq^2$
P2 $: q \times pq = pq^2$
P3 $: p \times q^2 = pq^2$

Discrete general covariance captures a kind of 'path independence' in the stochastic process that generates causal sets. It also implies that the probability of a particular causal set issuing from a growth process is independent of the order of birth we attribute to causal set elements. As noted, this is interpreted to mean that there is no physical significance to how we might order elements in the final causal set in terms of when they were added via the dynamics, since it is equally likely that the final causal set was birthed from multiple paths that, under a choice of gauge, involve different orderings of elements.

This brings us to Bell Causality. Bell Causality is a familiar notion of independence. The condition is added to capture the idea that elements of a causal set should only be influenced by their past, and not by other factors. In this way, the Bell Causality condition is close to what is sometimes called the causal Markov condition (not to be confused with the Markov sum rule above) which seeks to capture the idea that events are solely influenced by their immediate past.

To see the idea a bit more clearly, consider the following statement of Bell Causality from Rideout and Sorkin [26, p. 024002-6]:

> The ratio of the transitional probabilities leading to two possible children of a given causet depend only on the triad consisting of the two corresponding precursor sets and their union.

As noted, the precursor set for an element *e* in causal set C_{n+1} at stage *n* of the growth process is the set of elements that are earlier than *e* in C_{n+1}(and that are members of the causal set C_n at the previous stage of the growth process). What Bell Causality does is ensure that no elements of C_n other than those that are in the precursor set impact the probabilistic dependence of C_{n+1} on C_n.

Bell Causality can be stated more formally as follows:

$$\frac{P(C \to C_1)}{P(C \to C_2)} = \frac{P(B \to B_1)}{P(B \to B_2)}$$

In this equation, $C \to C_1$ represents the transition from a parent causal set to one of its children, as does the transition from $C \to C_2$. *B* is the union of the precursor sets for C_1 and C_2. $B \to B_1$ and $B \to B_2$ represent the transition from *B* to versions of C_1 and C_2 in which the same elements are added as in those sets, but any other elements have been pruned. Specifically, B_1 is *B* plus just the extra element added in C_1 and B_2 is *B* plus just the extra element added in C_2. The union of the precursor sets is

used because Bell Causality states a relationship between two children that may have different precursor sets, and so the two sets need to be combined for the relationship to be specified.

To see how Bell Causality works, consider Fig. 2. On the left-hand side, we have a parent causal set, C and two possible children C_1 and C_2. The precursor set of C_1 is just the bottom element in C, whereas the precursor set of C_2 is the first two elements in C. The union B of the two precursor sets thus contains the first two elements of C. On the right hand side, we start from B and then transition to causal sets containing just the elements added in the transitions from C to C_1 and C_2. Other elements of C_1 and C_2 have been removed from B_1 and B_2.

What Bell Causality tells us is that the elements that are missing on the right-hand side of Fig. 2 compared to the left-hand side make no difference to the transition probabilities between causal set C and its two children. That's what is captured by the equation above: the ratio of transition probabilities from B to B_1 and B_2 is the same as the ratio of transition probabilities from C to C_1 and C_2. This, in turn, captures the idea that the elements that are not in the past of new elements that are added during a transition between causal sets make no difference to the transition probabilities between those sets. Note that causal set C would have more than just these two children, and so in this sense Fig. 2 is an idealized picture, one that is simplified for the purposes of demonstrating Bell Causality. Bell Causality can be applied more generally to any number of children, by using it for every way of pairing a causal set's children.[7]

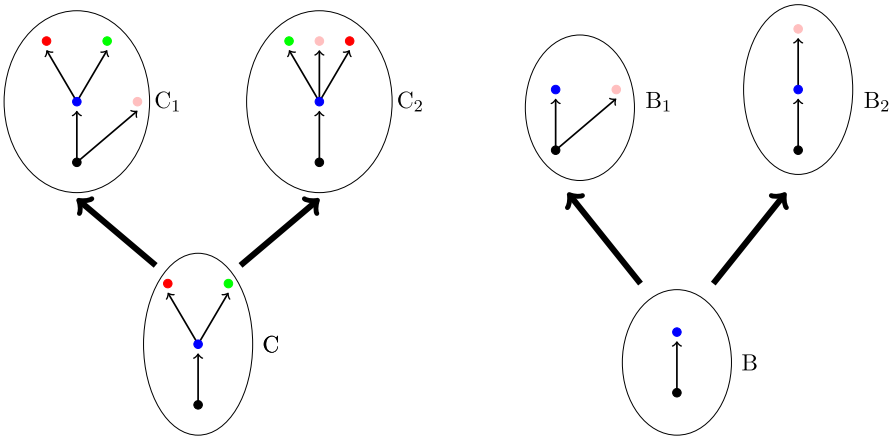

Fig. 2 C is the parent and C_1 and C_2 are children. B is the union of the precursor sets for elements added in C_1 and C_2, and B_1 and B_2 are causal sets formed by adding the new elements from C_1 and C_2 into B. According to Bell Causality, the transition probabilities on the left of Fig. 2 from parent to child, are the same as on the right of Fig. 2

[7] [37] interprets this to mean that events at a spacelike distance don't influence one another. This is reasonable, if we interpret the output of the growth process as something that is approximately isomorphic to a spatiotemporal manifold, with anti-chains corresponding to spacelike connections and chains corresponding to timelike connections. For then the Bell Causality condition would ensure that there is no non-local causation, which does appear to be necessary for recovering a spatiotemporal manifold.

4 Interventionism

Having outlined CST, we turn now to the interventionist approach to causation developed by [22] and [33], before applying it in §5. As noted, interventionism is here used to provide an account of actual causation—causation in the stronger sense alluded to in §2.

For present purposes, we will understand actual causation in terms of direct causation, which is defined as follows:

> [Direct Cause] A necessary and sufficient condition for X to be a (type-level) direct cause of Y with respect to a variable set $\mathbf{V}$ is that there be a possible intervention on X that will change Y or the probability distribution of Y when one holds fixed at some value all other variables Z_i in $\mathbf{V}$. [33, p. 59]

Direct causation is thus defined for variables, with respect to a variable set. A variable set is a way of modelling the different aspects of a given physical system. Woodward [33, p. 42] assumes that if X is a direct cause of Y, then X is a cause of Y, and so we will do the same.[8] We can read causation back into the physical system being modelled by inferring that direct causation between variables implies causation between the aspects of the system that they represent.

Each variable within the set can take a range of different values. The values of a variable represent possible ways for the different aspects of the physical system to be. A variable set is combined with a set of structural equations that link variables. The structural equations capture generalisations that relate the different aspects of a physical system.

So, for example, consider a very simple physical system involving: (i) the striking of a match, (ii) the match lighting, and (iii) the subsequent burning of someone's hand. For this system, we can introduce three variables: A, B and C. These variables represent the striking, the lighting and the burning respectively, and they constitute our variable set. In general, variables can have any number of distinct possible values. For this simple system, however, we can treat the variables as having binary values: 1 and 0, where 1 represents the occurrence of the event and 0 represents the

[8] The notion of a direct cause does not give Woodward's full picture of causation. For that, we need two further notions. First, the notion of a contributing cause:

> A necessary and sufficient condition for X to be a (type-level) *contributing cause* of Y with respect to variable set V is that (I) there be a directed path from X to Y such that each link in this path is a direct causal relationship; that is a set of variables $Z_1 \ldots Z_n$ such that X is a direct cause of Z_1, which in turn is a direct cause of Z_2, which is a direct cause of... Z_n, which is a direct cause of Y, and that (ii) there be some intervention on X that will change Y when all other variables in $\mathbf{V}$ that are not on this path are fixed at some value. If there is only one path P from X to Y or if the only alternative path from X to Y besides P contains no intermediate variables (i.e., is direct), then X is a contributing cause of Y as long as there is some intervention on X that will change the value of Y, for some values of the other variables in $\mathbf{V}$.

Next, the notion of a total cause:X is a total cause of Y if and only if there is a possible intervention on X that will change Y or the probability distribution of Y. [33, p.51]Together, Direct Cause, Contributing Cause and Total Cause "give us a way of fully capturing or cashing out the content of causal claims in terms of facts about what would happen under interventions" [33, p. 61].

non-occurrence of the event (e.g., A = 1 represents that the match is struck, A = 0 represents that the match is not struck).

This physical system, we can suppose, is governed by two generalisations: one that relates strikings to lightings, and one that relates lightings to burnings. These generalisations are encoded by structural equations that specify the way that the values of one variable are determined by the values of another. For this system, we have the very simple equations: B = A, C = B, which represent two facts: first, that the striking determines the lighting and, second, that the lighting determines the burning.

The set of variables and the set of structural equations can then be used to construct a graph. Each node in the graph corresponds to a variable, and links in the graph correspond to relationships specified by the structural equations between variables. If a relationship between variables X and Y is described by an equation, then X and Y are also linked in the graph. Together, the graph, the variables and the structural equations constitute a model. Note that within such a model, the variables that never appear on the left-hand side of a structural equation (or as the output of a probability function, see below) are commonly known as exogenous variables and the rest are endogenous variables. Figure 3 depicts a model for the simple match-striking system.

Once a model has been constructed, we can then use it to work out what causes what within the system. To do this we begin by setting the values of the variables in a model to their actual values (or, as we shall see below, by setting actual probability distributions over the values of some variables given the actual values of others). So, for the simple model above, we set the striking, the lighting and the burning to 1. Next, we consider interventions on the variables. For [33], an intervention on a variable is a possible cause that (i) changes the value of that variable from its actual value to some possible value and (ii) breaks the dependence of that variable on other variables via any structural equations. Importantly, interventions are not supposed to be anthropomorphic, in the sense that they are practically possible events that we, as agents, can bring about. The notion of an intervention transcends practical possibility (though by how much is something that we return to later on).

Woodward models interventions via the addition of an extra variable into a model, that acts as a 'switch' for the intervened upon variable such that the intervened upon variable is a function only of the intervention variable and nothing else. So, for instance, in the match-striking example, we can intervene to prevent the match from being struck, which would involve the addition of a new variable, I, and a new structural equation, $A = (1 - I)$, which effectively switches A off. Equally, however, we could add a new variable I and a new structural equation $B = (1 - I)$, and then remove the equation B = A. This leaves the striking in place, but cancels its effect on the lighting. In the first case, the intervention variable might correspond to restraining the person striking the match, in the second case, it might correspond to dousing the match with water.

By considering various interventions on the system, we can gather causal information. For instance, suppose we intervene on the striking to prevent it, to see what happens for the lighting, while holding the burning fixed. When we do this, we discover

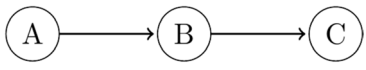

Fig. 3 A toy model. A is the striking, B is the lighting and C is the burning. The structural equations are: B = A, C = B

that the lighting does not happen (because of the structural equation that links the lighting to the striking). By the definition of a direct cause, we can thus conclude that the striking caused the lighting. Similarly, suppose we intervene on the striking to see its effect on the burning. To do this, we hold fixed all other variables. In this case, that means holding fixed the lighting. When we do this, we see that stopping the striking does not alter the burning, and so the striking is not a direct cause of the burning.[9]

In sum, then, the basic interventionist picture has three features: (i) a model consisting of a variable set, structural equations and a graph; (ii) an interventionist definition of direct causation and (iii) interventions that can be used to reveal causal information about a system. For our purposes, this basic interventionist picture needs to be expanded in two ways. The first concerns Woodward's notion of an intervention. This notion can be difficult to apply to physics [11]. To see this, suppose that we have a cosmological model of physics that describes a complete universe. Each variable corresponds to part of the universe, and the structural equations capture the lawful relations between parts. An intervention, in Woodward's sense, is the addition of a new possible cause that interrupts the system. But for a model of the universe as a whole, such a cause would need to come from *outside* the universe, which is difficult to understand.

A further issue concerns Woodward's specification of an intervention as a possible cause. In the case of causal set theory, we are interested in intervening on the elements of a causal set. Our goal is to use that framework to reveal causal dependence within the causal set structure. If, however, interventions require *adding causes* that effectively act as switches for other features of the system (inhibiting or activating those features), then in a sense we must presuppose that there is good sense to be made of interventionist causation within causal set theory already. Indeed, since the only things we can really add to a causal set are more elements, the only interventions we can make involve causal relations between causal set elements, whereby elements 'switch on' or 'switch off' other elements. The problem, then, is that working with Woodward's notion of an intervention would appear to be question-begging in the current context. It would require assuming, from the outset, that there is elemental causation operative in causal set theory.

To be clear, by raising these two problems we are not criticising Woodward's approach to interventionism. Nor do we rule it out that there is a way to make Woodward's approach work for our project or for applications to physics in general (indeed, it may work better in some cases, see below). The point is simply that Woodward's approach is not the right fit for what we are trying to do in this paper, and things work a bit more smoothly if we use a slightly different version of the interventionist framework. The way forward, for us, is to make use of the notion of a 'setting intervention'. A setting intervention is an intervention that involves simply setting some variable to a value "with no further restrictions on when such a setting operation is possible (or when it is permissible or legitimate to invoke it)" [35]. A setting intervention allows

[9] The striking is still a causal contributor to the burning. As noted in footnote 9, for [33], a variable X is a causal contributor to a variable Y when there is a chain of direct causes between X and Y. There is such a chain for the striking and the burning, because there is direct causation between the striking and the lighting and the lighting and the burning (for this second case note that an intervention on the lighting will switch off the burning).

one to change the value of a variable however one likes (so long as the resulting value is possible in the broadest sense), and then process the result. This is to be contrasted with the notion of an intervention specified in terms of possible causes (or possibility in general), which allows for only certain kinds of changes to the values of variables. Woodward [35] dubs these 'possibility constrained' interventions.

A setting intervention does not involve adding a new variable to a model. Rather, we simply change the values of variables directly, without the addition of a possible cause. This can be formally captured using Pearl's 'do-calculus', which involves the addition of an operation on variables, the 'do' operator. The 'do' operator is a function from the values of variables to other values. So, for instance, we can intervene on the striking by setting A to 0 (do (A = 0)) and we can intervene on the lighting by setting B to 0 (do (B = 0)).[10] As with possibility constrained interventions, setting interventions break structural equations. Figure 4 depicts the difference between possibility constrained and setting interventions on variable B for the match-lighting system.

Unlike possibility constrained interventions, setting interventions do not require the specification of possible causes external to the system. This serves our purposes well: a setting intervention does not require the possibility of anything 'impinging' on the system from the outside, and in this way we can make sense of causation in models of the universe as a whole. Moreover, a setting intervention is not analysed in terms of possible causes, and so we avoid begging important questions about causation in the context of CST.

We should note, however, that there is controversy over whether possibility constrained interventions or setting interventions should be used in applications of interventionism to physics. [24] and [11, 34] argues that the use of setting interventions is too liberal: it allows for the identification of causal structure where there shouldn't be any. For instance, a setting intervention can be used to identify causation between the matter and metric fields in general relativity, and between correlated particles in an EPR-type experiment. It is plausible, however, that these are cases in which causation is in fact absent. We take these examples seriously, but addressing them lies beyond the scope of this paper. For now, it is enough to say that we have sympathies for both

Fig. 4 Possibility constrained versus setting interventions. When circles are white, the value of the associated variable is 1, and 0 when black. On the left, a new variable is added to switch B off. On the right, the do operator sets B directly to 0 (do (B = 0)). In both cases the structural equation B=A is cancelled

[10] While we use Pearl's do-calculus to model setting interventions, Pearl does not endorse this picture of interventions. Indeed, Pearl's approach to interventionism is very similar to Woodward's. It is just as non-reductive about causation and Pearl seems to agree with Woodward insofar as he takes interventions that come from 'outside' of the universe to be problematic. Pearl makes this explicit when he writes "if you wish to include the whole universe in the model, causality disappears because interventions disappear" [22, p. 350].

Table 1 Probability distributions for values of B

	Light (l)	Not Light ($\neg l$)
Strike (s)	.9	.1
Not Strike ($\neg s$)	.1	.9

Table 2 Probability distributions for values of C

	Burning (r)	Not Burning ($\neg r$)
l	.9	.1
$\neg l$	.1	.9

sides of the debate, but that we ultimately side with Reutlinger and Frisch. Everything we say in what follows is thus contingent on the use of setting interventions.

The second expansion of the basic interventionist picture that we need is the addition of probabilities. We fit probabilities to a model by adding a probability distribution over the exogenous variables, $P(U)$, and by replacing the structural equations with a set of functions, f_i that fix the probabilities for endogenous variables $X_1 ... X_n$ based on the probabilities of exogenous variables $U_1 ... U_n$. Each such function has the form $X_i = f_i(\text{Parents } X_i, U_i)$. The parents of a variable are those variables that are directly connected to it. For instance, in the match-striking case, C has only one parent, B, and B has only one parent, A. Note that C is a descendent of A, but A is not C's parent. The set of functions gives us a probability distribution over the values of each endogenous variable for each possible value of parent variables and exogenous variables.

To see how this works, we can convert the matching-lighting case into a probabilistic example. To do that, we start by adding a probability distribution over the exogenous variables. In this case, only A (the striking) is exogenous. For simplicity, let us suppose that A has two possible values: striking and not striking, and that there's a good chance of the striking happening, say.9 (and thus.1 of the striking not happening). What we do is add a set of functions that set the probabilities of B given values of A, and the probabilities of C given values of A and B (Tables 1 and 2).

Starting with B, let us suppose that the lighting is very likely if the striking occurs (.9), but unlikely if the striking doesn't occur (.1). We can then map the probability distribution for B as follows:

C is more complicated, because we need to assign a probability distribution over its values for each combination of values for A and B. We can simplify matters, however, by applying the Markov Condition. For variables V representing nodes in a directed acyclic graph G, the Markov condition can be stated as follows:[11]

> For every variable X in V, and every set of variables $Y \subseteq V \backslash DE(X)$, $P(X|Ps(X)\&Y) = P(X|Ps(X))$

Where $V \backslash DE(X)$ is a set of variables in Y that excludes the descendants of X, and where Ps corresponds to the parents of X. What this condition does is effectively 'screen off' the probabilistic influence of A on C: the conditional probability of C on

[11] This is the 'screening off' condition stated by [13].

its parent B is the same as the conditional probability of C on B and A. We can thus specify probabilities and probability distributions for C in terms of B as follows:

The causal Markov condition is closely analogous to the Bell Causality condition imposed on causal sets (not the Markov sum rule). Both the Markov condition on interventionist models and the Bell Causality condition play a similar role, namely to induce a kind of probabilistic independence. Given the similar roles that Bell Causality and the Markov condition play in CST and interventionism respectively, we will impose the Markov condition on the interventionist framework in what follows. This is in line with the standard modelling practice for interventionism anyway, and ensures that the interventionist framework lines up with CST growth dynamics in an appropriate way by encoding something close to Bell Causality.

5 Application

We are now in a position to apply the interventionist framework to CST. We will do this in two stages. We proceed in this manner due to apparent disagreement over the extent to which the classical sequential growth dynamics for CST can be considered a real physical process. Thus, in the first instance, we will assume that the dynamics straightforwardly captures a real physical process, and apply interventionism under this assumption in §5.1.

Here we are largely following [9] and [29] who treat the growth process in this manner. Dowker notes that one can just run the dynamics to infinity, and then simply select 'completed' causal sets according to the probability measure given by the process, thereby effectively treating the growth process as a kind of fiction. However, she pushes back on this idea, writing that:

> ... it produces a different physical theory—some might want to call it a different interpretation of the theory—in which there is nothing to correlate with the occurrence of events. This blockified theory correlates less well with our sense experience than the one in which the birth of new spacetime atoms is a physical process. [9, p. 23]

[29] takes a similar line, noting that the growth process "offers us an active process of growth in which 'things really happen', but at the same time it honours general covariance". This interpretation of the growth dynamics is however controversial (for philosophical analysis, see, e.g., [1, 15, 38]). And so, in §5.2, we will consider this controversy in a bit of detail and then re-apply interventionism. As we'll see, even if we weaken the assumption that the growth dynamics corresponds to a real physical process, there remains scope for finding interventionist causation in CST. In §5.3, we consider objections, including the worry that the growth process should not be considered a physical process.

Why take the growth process as physical at all? We do this, in part, because (as can be seen from the quotes above) some physicists do treat it as a real physical process. Moreover, the idea that the growth process is physical has played some role in the

thinking around CST since the theory was developed, and so we believe it is important to consider this idea carefully.

However, our argument is not that just because physicists take the growth process seriously, that therefore we should too. Rather, we take this as an interesting starting point—an invitation—to see whether there is a way to take the process as physical, and what doing so would mean exactly. Our analysis reveals that it likely involves treating the process as a causal process, whereby causal set elements actually cause one another to exist in the sense of elemental causation. Thus, as noted in §2, rather than thinking of the growth process as passage or becoming, we can think of it as causation. In this way taking the growth process as real is worthwhile for our analysis, since doing so reveals an alternative way to understand this process, which is significant given the troubles with the current way of interpreting it as physical.

5.1 Stage One

We start with an observation. Consider the toy model of sequential growth dynamics from Fig. 1, which depicts the growth of a causal set up to three elements (noting, again, that this does not represent the entire, potentially infinite growth process). Now, consider again the following causal set:

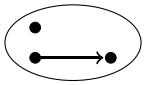

Call a causal set of this type: a Type 1 set. As can be seen from Fig. 1, there are three causal sets of this type. The transition probabilities toward this single physical state appear to depend on what happens at a previous stage of the growth process. If, at the second stage of the growth process, an element is added that is unlinked to any existing elements, then the transition probability to each Type 1 set is pq. If, however, at the second stage, an element is added that is linked to existing elements, then the transition probability to each Type 1 is q^2. Thus, since generally $pq \neq q^2$, it follows that the transition probabilities to Type 1 causal sets at the third stage differ, depending on what happens at the second stage.[12]

[12] One might worry that the transition probabilities only diverge when $p \neq q$. That is true for this particular case (though ultimately, it doesn't matter for the application of interventionism, see below). However, it is not true in general. For instance, consider the following four element causal set:

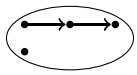

Call this a Type 3 set. There are four paths to this causal set. Three of them issue from the Type 1 causal set discussed above. The fourth path, by contrast, issues from the following causal set:

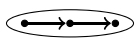

Call this a Type 2 set. For each transition from a Type 1 set to a Type 3 set, the transition probability is the same: pq. However, for the transition from a Type 2 set to a Type 3 set the transition probability is q^3(by the transition probability rule with $n = 3$, $m = 0$, and $\varpi = 0$). Note that $pq \neq q^3$ even when $p = q$.

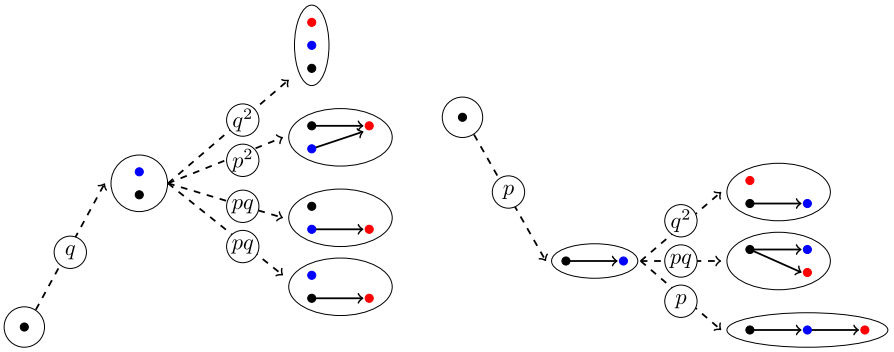

Fig. 5 A switching effect in the growth dynamics. At the second stage, the growth process is switched, altering the transition probabilities

Fig. 6 Intervening on the growth process. White nodes have value 1, black nodes have value 0. We intervene on B with a setting intervention modelled by the do-calculus: do (B = 0)

Table 3 Probability distribution for values of A

a_1	a_2
p	$q = (1 - p)$

We can thus imagine 'switching' the growth process at the second stage from one causal set to another, resulting in a change in the available transition probabilities to a Type 1 causal set at the third stage (Fig. 5). This can be depicted as follows:

This switching is indicative of causation. To show this, we can capture the switching within an interventionist model. To do this, we will assume that the first element of the causal set has been generated, and look to intervene on the addition of the second and third elements. We thus add two binary variables: A with possible values a_1 and a_2 and B with possible values b_1 and b_2. The values for A correspond to the two possible causal set types that can exist at the second stage of the growth process. The values for B, by contrast, correspond either to a situation in which a Type 1 set is produced at the third stage or to a situation in which some causal set of another type from the available options is produced at that stage. A and B represent nodes in a simple two-node, directed acyclic graph, such that A is the parent of B (Fig. 6).

Next, we add an initial probability distribution over the values of A: a_1 and a_2. Since a_1 is the addition of an element that is linked to existing elements, and a_2 is the addition of an element that is unlinked to existing elements, and because we know from the dynamics that the addition of a linked element at this stage occurs with probability p and the addition of an unlinked element occurs with probability $q = (1 - p)$, we have the following initial probability distribution over a_1 and a_2:

Now, we need a probability distribution over values of B in terms of values of A. We have to be a bit careful here. We can't let any probability distribution for values of B correspond to net transition probabilities for the generation of Type 1 sets (since

these are all the same, regardless of the path taken through the dynamics). Nor can we let the probability distribution correspond straightforwardly to the transition probabilities for causal sets into the third stage, since the transition probabilities are for labelled causal sets—which are tokens—and the values for *B* correspond to types (namely the presence/absence of a Type 1 set at the third stage).

One natural way forward is to derive a probability distribution for values of *B* from transition probabilities as follows. Suppose that at the second stage a coloured causal set of two unlinked elements is generated. This is the situation corresponding to a_2. We know by the dynamics that there are four coloured causal sets that can be produced at the third stage. Two of these are of the same type (Type 1). Thus, to work out the overall probability of there being a Type 1 causal set at the third stage, we can add the transition probabilities for the two ways of producing that set. Thus, because each way of producing a Type 1 causal set has a transition probability of *pq*, we can set an overall probability of $2pq$ for b_1 given a_1. Since by the Markov sum rule all transition probabilities must sum to 1, we can also use $1 - 2pq$ as the probability for b_2 given a_1.

The same reasoning can be applied to the situation in which a causal set of two linked elements is generated at the second stage, i.e. a_1. In this case, however, there is just one way of producing a Type 1 causal set with transition probability q^2, and so we can use the transition probability q^2 for b_1 given a_2. As before, using the Markov sum rule, we let $1 - q^2$ be the probability of b_2 given a_2. Thus, we have the following probability distribution over values for B in terms of values for A:

We now have a precise way to consider interventions into the toy model of growth dynamics (Table 4). Suppose that, as a matter of fact, the system moves from the first stage to the second stage. We thus set the value of A to a_1, meaning that the system adds a second element, and that element is connected to the first element by a link. This determines the probability distribution at *B* as the second row in Table 4.

Next, we imagine an intervention into the system at the second stage whereby we force the system to take an alternative path. Thus, instead of a_1, we intervene to switch the system to a_2 at the second stage (i.e. do (A = a_2)). This, in turn, changes the probability distribution over the options at B. In particular, the probability distribution over the options at the third stage switches from the second row in Table 4 to the first row. We can expect the two probability distributions to be different because $2pq \neq q^2$ generally speaking. This hypothetical intervention therefore satisfies Woodward's definition of a direct cause specified in §3.1. For recall that, according to that definition, it is sufficient for a variable X to be a direct cause of a variable Y if intervening on X alters the probability distribution at Y. That is what happens in our model: when we intervene on A the probability distribution over B changes.

Is the probabilistic dependence of the causal set at the third stage on the causal set at the second stage indicative of elemental causation? We believe so. For the two scenarios being 'switched' between are quite different. In one scenario, the third element that is added at the third stage (the red element) is linked to existing elements, in the

Table 4 Probability distributions for values of B

	b_1	b_2
a_1	2pq	$1 - 2pq$
a_2	q^2	$1 - q^2$

second scenario it is unlinked. But it is plausible that causal set elements are individuated by the relations they stand in [37]. Thus, what we are seeing is two scenarios in which different elements are added, depending on what happens at the second stage. Our analysis thus reveals that the existence of a particular element depends, with a certain probability, on which elements exist at the second stage.

There are two points to note about this result. First, the result is based on a choice of variables. A different choice of variables might have a different outcome. For instance, if one lets the values of B correspond to each possible causal set at the third stage, then it is unlikely that the shift in probabilities that we have identified will show up. That shift relies on taking B to correspond either to the generation of a Type 1 set or to its failure to generate. We have built the model so that the probability distribution applies in this way. It is tempting to think that, as a result, the causation we have identified is somehow 'subjective', based on a choice that we make. But that's not the case: so long as there is some choice of variables that reveals causal dependency and that satisfies the requirements of both interventionism and the dynamics of the system at hand, we have found a causal connection. It is just that for other ways of representing the same system using variables, this causal connection may not be discoverable using interventionism. But it doesn't follow that the connection is not a part of the system. The choice of variables is a way of finding causal connections in a system by describing it in the right way; it is not a way of inventing causal connections through one's description.

Second, we can expect a similar 'switching' intervention to be applicable throughout the growth dynamics. There are many causal set types that can be reached via multiple paths. Discrete general covariance tells us that the net transition probabilities toward any such set will be the same. However, it is generally compatible with this that the transition probabilities that lead immediately into a causal set of a given type from the previous stage are not all the same (keeping in mind that the transition probabilities and the net transition probabilities are different things). That is, for a causal set of type *T* that arises at a stage *n* of the growth dynamics, it is not in general the case that the transition probabilities from causal sets at $n-1$ to *T* at *n* are all identical. Quite the opposite: they often differ.[13]

This provides a basis for the switching behaviour we have identified. For a causal set of type *T* that arises at a stage *n* of the growth dynamics, if there are at least two causal sets C_1 and C_2 at stage $n-1$ such that the transition probability from C_1 at $n-1$ to *T* at *n* is not the same as the transition probability from C_2 at $n-1$ to *T* at *n*, then intervening to 'switch' the system from C_1 to C_2 at $n-1$ will influence the transition probability from $n-1$ to *T* at *n*. This switching can generally be captured by an interventionist model of two binary variables *A* and *B*, where *A* corresponds to the choice between C_1 and C_2; *B* corresponds to the occurrence or non-occurrence of a set of type *T* and where the probability distribution over values of *B* is determined

[13] Indeed, so far as we can tell, when, for a causal set of type *T* at *n* there are two causal sets C_1 and C_2 of different types T_1 and T_2 at $n-1$ such that there is a transition probability from those causal sets to a *T*-type set at *n*, those transition probabilities will generally diverge, even while the net transition probabilities agree.

by the transition probabilities from C_1/C_2 into T using the same method as above.[14] Thus, we can in general model interventions that correspond to the switching we have identified, and these cases will satisfy the notion of a direct cause. There is thus some reason to think that there will be actual causation throughout the growth dynamics. We've already suggested that this actual causation is elemental causation. We will strengthen this conviction in §5.2.

5.2 Stage Two

In the previous section, we assumed that the sequential growth dynamics corresponds to a real physical process. As noted, we have proceeded in this way in the first instance, as it seems to align with the way that some causal set theorists think about the view. However, there are reasons to doubt this picture of CST.

The problem comes this way: discrete general covariance seems to inject a peculiar indeterminacy into the growth dynamics. To see this, consider again the toy model in Fig. 1, and consider a Type 1 set. As we've seen, there are three paths to a causal set of this type, and the net transition probabilities along each path are the same. As also discussed, this is interpreted to mean that there is no physical difference between the three coloured versions of a Type 1 set in Fig. 1. More than this, however, it is taken to imply that the labelling used to produce the causal sets is not physically meaningful. The labelling is a mere choice of gauge. There is thus no fact of the matter as to which causal set elements came into existence 'first' beyond the constraint imposed by internal temporality.

This has striking implications. If we look at the three paths that produce a Type 1 set, we see that they are quite different at the second stage. On some paths, there are two unlinked causal set elements at the second stage, whereas on other paths there are two linked causal set elements at the second stage. As [38] note, there seems to be no fact of the matter as to which causal set exists at the second stage. The problem generalises in an uncomfortable way: there appears to be no fact of the matter as to what exists at the *nth* stage for any finite stopping point, because the probability measure of causal sets is only well-defined in the limit. It is thus only once the dynamics is run to infinity that the causal set 'snaps' into place.

It is tempting to therefore conclude that the *only* real causal set is the one produced in the limit, and that the causal sets produced along the way are not physically real. If the causal sets produced along the way are not physically real, then there is no sense to be made of intervening at a given stage in the growth process to switch it to a different causal set.[15]

As[38, p. 922] note, however, discrete general covariance does not completely undermine the prospects for interpreting the sequential growth dynamics as a real physical process. This is so for two reasons. First, they note that many causal set

[14] The probability distribution over A is a bit tricky. Having tested this on a range of different causal sets, it is clear that an idealisation must be invoked: namely, we must assume that the values of A are exhaustive, and thus that the probability distribution over its values is essentially that given in Table 3.

[15] As Wüthrich and Callender (2017, p. 920) remark, one potential way around this is to develop a 'hidden variables' interpretation of the causal set dynamics. [10] considers this, but it is criticised by [1]. Such an interpretation, if viable, would provide a basis for the interventionist analysis in §4.1.

models feature 'posts'. A post is an element that is 'comparable' to all other elements in the causal set. That it is comparable means that it's either later or earlier than all other elements (no elements are disconnected from it). A causal set model can reach a post in a finite number of steps. When it does, everything earlier than the post snaps into determinate existence, despite the growth process being incomplete (insofar as it has not yet reached the infinite limit). This suggests that we can make sense of the growth dynamics reaching some definite, determined stopping points along the way toward future-infinite completeness.

Second, [38, p. 923] point out that the cardinality of the causal set is generally well-defined for each stage of the growth process, even if the process itself is only completed when run to infinity. Thus, while we cannot say exactly which causal set exists at the *nth* stage of the growth process for any finite *n*, we can say what the cardinality of the causal set at that stage will be, namely: *n*. Moreover, we also know that the cardinality strictly increases with the procession of the growth dynamics. Thus, the cardinality of the causal set for any stage of the growth process *n* is strictly smaller than the cardinality for any stage of the growth process $n + 1$.

Do these two points help us? Take the presence of posts first. At first glance, a post presents a definite stopping point, and so offers a point at which one might apply interventions. However, a first problem is that whether any element is a post is only specified in the limit [38, pp. 922-923]. At any finite point in the growth process, for any element *e* such that every element is ordered with respect to *e*, there can be a new element added later on that is not ordered with respect to *e*. Thus, we seem to face the same problem as above: it is unclear that there is after all a determined stopping point in the growth process that is open to intervention. A second issue is that it's not clear that it's enough for our purposes to show stages involving posts are stages of real physical processes. It would be better if the interventionist machinery could capture causal relations in causets in general, and not in specific stages involving posts only.

The second point made by [38] provides a better basis for our analysis. One important aspect of the causal set dynamics is that it is *step-wise*. This means that for a causal set to grow from *n* elements to $n + 2$ elements, it must first grow to $n + 1$ elements. So, for instance, there is no way to 'jump' from a causal set of 2 elements to a causal set of 4 elements. Such transitions go beyond the standard sequential growth dynamics insofar as there is no dynamical mechanism by which causal sets can grow by more than one element at a time.

Given this, we can consider a different kind of intervention into the dynamics. Rather than intervening to 'switch' the process from one causal set to another at a stage, we intervene to 'prevent' the growth of an element at a stage. Intuitively, preventing the addition of any element at a stage *n* of the growth process will prevent the addition of any element at stage $n + 1$. That's because if we prevent the addition of an element at stage *n*, then for the growth process to continue, there must be growth from a causal set of $n - 1$ elements to a causal set of $n + 1$ elements. But that is dynamically disallowed. If we intervene to prevent the growth of an element, we effectively intervene to prematurely halt the dynamical process.

We can easily build an interventionist model to reflect this kind of change. To do this, we take toy the model of the dynamics in Fig. 1 and introduce a set of three variables: A, B and C. Each variable is binary, and represents either an element being

added (1), or no element being added (0). Values of A represent the first element being added or not being added; values of B represent the second element being added or not being added; values of C represent the third element being added or not being added.

The structure of the model is just this: the second element is added just in case the first element is added and the third element is added just in case the second element is added. We thus have two structural equations for the system: $B = A$ and $C = B$. We can now consider interventions on the system to prevent the growth of causal sets. Thus, we can imagine intervening on *B* to prevent the addition of the second element. When we do this, we thereby prevent the addition of the third element. We can model this intervention as follows:

We can easily extend this model for any finite stage in the growth process. To do that we introduce variables $v_1 ... v_n$ where the *ith* variable represents the addition or non-addition of the *nth* element. We then add a set of equations, *E*, where each equation says that the value of the $i + 1th$ variable is equal to the value of the *ith* variable, for $i > 1$. For any sized model, if we intervene on the *ith* variable to switch it from 1 to 0, this will set the $i + 1th$ variable to 0 as well, which will have a knock-on effect for all other variables $j > i + 1$. In other words, if we halt the addition of an element at any stage *n*, then we halt the addition of an element at the $n + 1th$ stage and for all stages $m > n$.[16]

What this shows is that if for any finite stage there is a determinate cardinality for the causal set that exists, then there is also causal dependence between each stage of the growth process. In particular, there is a model in which the addition of the *nth* element is a direct cause of the addition of the $n + 1th$ element, which is indicative of elemental causation.

Note that this is different to the weak sense of causation identified in §2, involving causal connectibility. Recall that this weak notion of causation is present in general relativity. The notion of causation we have identified in this section goes beyond that notion of causation. For it is not the case, so far as we know, that there is elemental causation in general relativity: spacetime points don't actually cause one another to exist in any substantive sense (more on this in §5.3). That is, if we imagine an intervention that removes a spacetime point, it is far from clear that this will wipe out any other spacetime points. In the case of CST, however, it does seem that intervening to wipe out a causal set element will wipe out elements further along in the growth process. If that's right, though, then even the fairly minimal form of dependence between the stages of the growth process identified here is enough to establish that CST is causal in the strong sense of using actual, elemental causation in the dynamics.

[16] Indeed, in principle we can extend this to produce an infinitary interventionist model that captures the entire growth process. For this, we let the variable set be infinite, thereby allowing one variable for every step in the growth process all the way into the final, infinite causal set. When we do this, we can see that for any variable v_n there is an intervention on that variable that alters the next variable in the chain v_{n+1}.

5.3 Objections

We anticipate five objections. First, one might argue that the dependence of causal set elements on one another is already implicit in the sequential growth dynamics. The appeal to interventionism doesn't really add anything. We disagree. What's perhaps implicit in the growth dynamics is that there's dependence of some kind between causal set elements as they are added through the growth process. What the application of interventionism shows is that the dependence at issue corresponds to a case of elemental causation. Thus, our analysis tells us what type of dependence is implicit in the growth dynamics. As discussed in §2, this is a useful thing to know about CST.

One might demur: interventionism is a flexible framework. Too flexible, one might argue, to reasonably conclude that the kind of dependence identified is causal dependence. In particular, interventionism, one might argue, can be applied to cases of grounding. That is, we can consider interventions on the grounds that make a difference to the grounded. That being so, one might worry that the application of interventionism in this paper underdetermines the kind of dependence identified in CST. Perhaps it is causation, but also it might be grounding.

If the dependence is grounding, that's still significant. It still helps us to better understand the way that CST works, and the implications of taking the growth process to be a real, physical process. It also helps us to understand the way that spacetime emergence might work, suggesting that a grounding-based account or similar might be appropriate. Of course, the issue here is that the use of interventionism might underdetermine whether it is causal or grounding dependence. Nonetheless, even if what we find is that there's a metaphysically robust kind of dependence at work in the dynamics of CST, that's still worth knowing.

Ultimately, however, we think that the type of dependence is not underdetermined. Those who take interventionism to be applicable to grounding also generally regard this as a reason for collapsing the distinction between causation and grounding to some degree, treating both kinds of dependence as a kind of causation [28, 31]. In that case, our application of interventionism reveals causation after all. If one is not inclined to collapse the distinction, then this is usually due to perceived differences between causation and grounding. But these differences can be used to break the underdetermination here as well.

Here we have in mind one prominent account of the difference between grounding and causation [32].[17] On this account, the difference is this: causal relations are governed by the physical laws of nature, whereas grounding relations are governed by modally broader constraints, sometimes called 'metaphysical laws'. The interventions we have considered are based in the dynamics of CST, and in this sense are responsive to physical laws laid down by the theory. There is no sense that we can see in which any broader metaphysical laws are coming into play. We thus have

[17] This is not the only account. Another account appeals to time. Roughly: grounding happens at a time, whereas causation happens across time. This account is generally thought to fail, though see [3] for an attempt to revive it. If this account is assumed, and the growth process is considered temporal as some proponents of CST maintain, then this would be another way to break the underdetermination.

good reason to suppose that the relations being identified are causal relations, and not grounding relations.[18]

Second objection: causal connectibility of the kind introduced in §2 just implies the presence of interventionist causation. That being so, it is quite straightforward that if there is causal connectibility (which in the case of CST, we can suppose there is) that then there is interventionist causation. But this also makes interventionist causation seem too easy to obtain, and thus makes our results trivial.

As briefly noted in §2, however, the presence of connectibility is not sufficient for the presence of actual causation. If there is causal connectibility between *x* and *y* then all that tells us is that it's possible for some actual causal relation to be present between *x* and *y*. What it doesn't tell us is whether, in fact, were one to make a change to *x* in the interventionist sense, there would be a change in *y*. It is entirely compatible with the presence of causal connectibility that, in fact, there is no change in *x* that would make a difference to *y*.

The fact that interventionism analyses causation modally makes no difference to this, for the kind of modal link provided by causal connectibility is different to the one provided by interventionism. Given the presence of causal connectibility, it is *possible* that, were a change made at x, there would be a change at y. This is different, however, to it being actually the case that, were a change made at x, there would be a change at y. The difference here can be revealed in the logic of the two notions. If we let '$A \Box\!\!\rightarrow B$' stand for an interventionist counterfactual (an intervention on A makes a difference to B), then what causal connectibility tells us is that $\diamond(A \Box\!\!\rightarrow B)$. This is compatible with $A \Box\!\!\rightarrow B$ being actually false, in the way that $\diamond A$ does not generally imply A. What interventionism adds, is the claim that the counterfactual is actually true.

Third objection: it is trivial that there is actual causation in physics, given our analysis. The objection can be framed using general relativity. Suppose we take two spacetime points *x* and *y*, and add a matter field over those points, along with dynamical equations concerning the way in which alterations to the matter field at *x* result in alterations to the matter field at *y*. Then surely this now is sufficient for actual causation. But if that's right, then it's straightforward to find actual causation whenever there's dynamics of this type.

There are a couple of things to say here. First, there is serious disagreement about whether interventionism works in all such cases. Indeed, in the case of general relativity, counterfactual theories such as interventionism will not work [17]. We lack the space to go into the details here, but the basic idea is that the dynamical information fails to support the strong 'would' of the counterfactual. Thus, even though there are dynamical equations linking matter fields at points, this does not always result in causation. It really depends on whether the dynamics supports interventionist counterfactuals, which is something that generally needs to be shown on a case-by-case basis.

[18] As we note below, thinking through the intervention requires breaking the physical laws. But this is often true for interventions for causation. It doesn't mean that the relations are not causal relations, since it's compatible with the imagined intervention breaking the laws that the actual relation is governed by those laws.

Second, there's an important difference to be drawn between interventions on matter fields located at spacetime points, versus interventions *on the points themselves*. That is, we might have dynamical equations in general relativity, and many other theories, which support interventions on matter at a point, with this ramifying through that point's forward light-cone (though the just-cited disagreement about general relativity means even this may not be straightforward). But even if we have that, we don't seem to have interventions on the points themselves. The point can be put in terms of our distinction between elemental and material causation. That material causation is possible does not imply that there is any elemental causation.

We thus need to be very careful to distinguish two kinds of interventions in general relativity: an intervention on some A located at a point x, versus an intervention on x itself. Likewise, we should distinguish two sorts of potential interventions in CST: on the CST elements, versus on things located at those elements. It is the first kind of intervention that we are considering here as no matter fields are available to perform the second kind of intervention.

One might respond, however, that the dynamical equations of a theory can link just the points, and indeed that such dynamical equations do link points in general relativity. For instance, in vacuum solutions it looks as though we have dynamical equations that relate spacetime points in the absence of matter.

But while this may be so, this is still not enough for elemental causation. To illustrate this point, imagine that we perform an intervention on a general relativistic spacetime to 'remove' a spacetime point from the manifold entirely. This is analogous to the intervention on a point in CST to remove an element at a stage of the growth process. In the case of general relativity, it is *not* the case that, were we to remove a spacetime point, that every point in its forward light cone would be 'wiped out'. Indeed, this would conflict with the way that the geometry of spacetime is determined locally by the energy-momentum tensor. Instead, it seems more likely that a singular spacetime would be created (roughly a spacetime with a 'hole' in it).

Thus, the case of general relativity shows us that having some points and some dynamics is not enough to get us interventions that support elemental causation. The dynamics needs to induce a dependence between points, such that wiping out one forces the removal of others. This requires a special *kind* of dynamics, namely the type of dynamics that we have when we take the growth process in CST physically seriously.

Thus, even granting that dynamical equations can give us interventionist causation (which, as noted, it is not obvious we should), they may not give us elemental causation, which is the kind of causation we are focusing on. This type of causation is unusual, and reveals a type of causal dependence between the elements quite unlike what we find in other theories, including general relativity.

Of course, one could add a growth process into a theory. For instance, one could take general relativity and add a physical growth process whereby spacetime points dynamically depend for their existence on one another through relations of actual causation. If one were to do this, then general relativity might admit of the same analysis as we have applied to CST, and so perhaps could be considered causal in the sense we have in mind (involving elemental causation between the points of space-

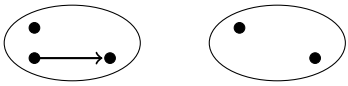

Fig. 7 Intervening on the left-hand causal set to produce the one on the right

time). But, as noted, this doesn't seem to be a feature of general relativity as it stands, given the way the geometry is defined locally at a point.

Fourth objection: one might worry that the growth process is not to be treated as physical in any sense, even the minimal sense that we have described here. To be clear, we are not defending the idea that the growth process is physical. Rather, we are assuming it is for the sake of argument, and then demonstrating the way that the growth process so understood yields actual, elemental causation. In this respect, our approach is similar to the one taken by Wüthrich and Callender [38], who also take the growth process to be physical and aim to interpret CST under this assumption. Our conclusions, like theirs, are beholden to this assumption. We think the assumption is worth making, however, since some of the physicists who developed CST seem to interpret the theory in this way, taking the growth process as a physical aspect of the theory. We take the perspective of physicists seriously in this respect, and consider the implications of this interpretation of the theory.

As discussed, doing so opens up a new way to interpret that process as a real physical process. Rather than interpreting the process as one of temporal passage or becoming, we consider it as a causal process. As mentioned in §2, this is quite useful. For while indeterminacy is not considered a feature of passage, it does seem to be a feature of causation in general, and so the worries that have been raised against taking the growth process as physical appear to have less force. Though, of course, whether this new interpretation is problem-free, remains to be seen. That is not an issue we have space to deal with here, but hope to examine it in future work.

Still, one might claim that the only causal sets that we should take seriously are the ones produced in the limit, and thus it is only infinite 'complete' causal sets that exist. Even the stages of the growth process with their increasing cardinality should not be taken physically seriously. We admit that under this blockhead interpretation of CST, the application of interventionism we have identified does not seem to work. To see this, suppose we have just a completed causal set and consider 'removing' an element with an intervention in the way that we considered for the weakened conception of the growth process in §5.2. For simplicity, suppose we just imagine intervening to remove an element from a three-element causal set. There are two options for what might follow from making this change. One option is depicted in Fig. 7.

In this case, we have removed the bottom left element in the left-hand causal set. Doing so, however, does not wipe out the element ahead of it in the causal set ordering. Instead, the result is just a new anti-chain in the causal set on the right. A second option is that the causal set on the left is turned into a one-element causal set, since when we remove the bottom left element in the left-hand causal set we also wipe out every element that is later than it in the causal set ordering.

The difficulty is that without relying on the dynamics, there is no basis for preferring one of these outcomes of the intervention to the other. Both outcomes result in possible causal sets, and neither is privileged from the perspective of what would happen were we to remove an element.

More generally, without the dynamics, there is no obvious sense in which wiping out a single element results in the elimination of all causal set elements 'ahead' of the removed element in the causal set ordering. When we remove an element, one possible outcome is the loss of everything ahead of it, but another possible outcome is the loss of just that element, with everything else remaining intact. Again, both options are possible causal sets. Without a way to select between these, we cannot say that removing an element would have any particular causal outcome for other elements. Thus, if we set aside the growth process and its dynamics, there doesn't seem to be a way to find causal dependence between causal set elements of the kind needed for elemental causation.

If that's right, then perhaps our conclusion can be strengthened. We have argued that if the growth process is a real physical process, then causal set theory features elemental causation. But given what we've just said, it may be that causal set theory features elemental causation just in case the growth process is a real physical process. Then elemental causation is only really a core part of the theory under a specific interpretation.

Whether this stronger conclusion is justified remains to be seen. What we need to do is take a closer look at the 'kinematics' of a completed causal set and see what happens when interventions are applied there.[19] It may be, for instance, that 'deleting' elements of the causal set through interventions does have a knock-on effect throughout the rest of the causal set, despite what we've just said. For it may be that the dynamics imputes dependence between causal set elements, even if it is not considered a real physical process in the manner described in §5.2. Perhaps if we rewrite the dynamics in a way that doesn't rely on stages, and then use this revamped dynamics as the basis for an interventionist analysis, we'll uncover elemental causation once more. We note this as a potential avenue for future work.

Fifth objection: one might note that there is no possible cause in CST that can prevent the growth of a causal set element. Thus, there can be no intervention. Recall, however, that we are not using a notion of intervention that is tied to possible causes. Still, one might worry that even if an intervention does not correspond to a possible cause, it should still correspond to a physical possibility *in some sense*. The concern, then, is that there is no dynamical procedure at all for preventing the growth of the causal set, and so the case we are considering is one that is simply not physically possible by the lights of CST. But, again, there is no requirement that setting interventions should be constrained by any particular notion of possibility. So the setting intervention does not need to be underpinned by a possible dynamical procedure described by the growth dynamics.

One might *still* be worried. Even a setting intervention, it might be thought, should correspond to an operation on the models of a theory, in this sense: a setting intervention should correspond to an operation that takes one from a model M of a theory to a model $M* \neq M$ of the same theory. One might argue, however, that the interven-

[19] We are hesitant to use the word kinematics here, as the distinction between kinematics and dynamics could lose its meaning in this context. If there is no physical process, but only a final complete causal set, then the dynamical laws simply express the internal organisation of the causal set—just as in a Humean approach to spatiotemporal theories, laws are seen as describing spatiotemporal regularities in the distribution of events.

tions we are considering don't amount to an operation on models of this kind. That's because the intervention we have in mind is an intervention into an infinite growth process that halts it at some finite step, leaving us with a finite causal set. If the only dynamical models of causal set theory are ones where the causal set grows to infinite size, then our intervention does not correspond to an operation on models in the relevant sense.

There are two things to say here. First, there are finite causal sets in the kinematics of causal set theory. So it is not clear that the intervention at issue takes us beyond the theory in any objectionable sense. Granted, it takes us from an infinite causal set to a finite one, but both are, in some sense, part of the broader ontology of the theory, and so we think there is still good sense to be made of the intervention. Second, it is far from clear that interventions ought to be operations on models in the relevant sense. It certainly seems that the use of the interventionist framework must be much broader than this kind of operation already.

For instance, suppose that we have a set of deterministic dynamical laws and we want to know what happens if, under those laws and for a specific set of initial conditions, we intervene at a particular time t to prevent an event e from happening. Any such intervention takes us beyond the model space of the theory, since there is no model that leads from the relevant initial conditions to a situation in which e is missing via the deterministic laws. This kind of case, however, seems like a very natural use of the interventionist framework. It roughly corresponds to one in which a small 'miracle' wipes out the event e.[20]

As [35] puts the point when considering the counterfactuals that correspond to what would happen under interventions:

> ... what is crucial is not whether the antecedent of the relevant counterfactual is nomologically or physically possible but rather whether we are in possession of well-grounded scientific theories and accompanying mathematics that allow us to reliably answer questions about what would happen under the supposition of such antecedents. We count interventions as "possible" as long as this is the case.

The interventions we have considered for CST appear to be possible in the sense that Woodward indicates. They are based on well-grounded scientific theories and accompanying mathematics that allow us to answer questions about what would happen were causal set growth to be prevented. One might challenge the idea that CST is well-grounded, noting that it is not empirically confirmed. But that is surely too narrow a definition of 'well-grounded'. We take it that a reasonable understanding of 'well-grounded' would include sufficiently well-developed and serious scientific theories and associated mathematical frameworks. CST surely qualifies as well-grounded in this sense.

[20] There is, of course, a delicate question concerning how, on the one hand, we can hold the dynamics fixed so that we can reason through the implications of an intervention while, on the other hand, violating the dynamics to a certain extent. This is a general issue with counterfactuals of this kind, however, and not one that we can hope to resolve here.

One might disagree: it has not yet been shown that the dynamics of causal set theory will produce, with high probability, manifoldlike causal sets that approximate spacetimes of the type found in general relativity. This is a fair point, so we ask for some latitude in what 'well-grounded' means for the application of interventionism for philosophical purposes. As we see it, the notion of 'well-grounded' is supposed to capture, in part, the idea that a theory is sufficiently well-developed that the application of interventionism is coherent, and provides determinate answers. That is indeed the case, as our analysis shows. But this seems to be enough for our purposes, because we are drawing conclusions only about causal set theory itself. Matters would be different if we were trying to draw conclusions about the world, or about another theory beyond CST. But that's not what we're doing.

One might disagree, again, arguing that CST plus the growth process construed as physical is not a well-grounded scientific theory in even the weak sense just outlined: a theory that is sufficiently well-developed that it can admit of philosophical analysis. In reply, we suggest that if a theory is taken seriously by physicists, then philosophers have some reason to take it to be well-grounded in this sense, using it as a reasonable starting point for philosophical analysis. As discussed, some physicists do take seriously the combination of CST and the claim that the growth process is a real physical process.

Again, just to clarify, our argument here is not a straightforward argument from authority. We are not simply saying that because physicists take the growth process to be physical, we should too. Our argument is that this is a reasonable starting point for our analysis, and so a reasonable starting point for applying interventionism to CST to reveal elemental causation. Indeed, given our analysis, it may be that we have more reason to take the growth process physically seriously, since we now have another way to interpret it. In this way, the analysis justifies itself: by performing the interventionist analysis on CST, we can provide more support for taking the growth process physically seriously. Of course, it could be still that the process should not be taken seriously in this way. But that is a further question.

We admit that if the growth process were shown to be contradictory, then that would offer a defeater against what we do in this paper, but this has not been shown so far. One might respond that the growing block theory of time has been shown to be contradictory, and it is sometimes thought that CST operates just like a growing block theory. But no one has yet shown that the growth process in CST is just like the growing block theory. At best, the two processes have been shown to be somewhat analogous. Moreover, as we've been at pains to point out, what we do here shows a way to interpret CST without using a growing block or becoming model. That being so, we don't yet see a reason to doubt that CST is a well-grounded scientific theory, even under the interpretation that the growth process is physical in some sense. More to the point: our analysis steers CST away from a temporal, growing block interpretation. Rather than viewing the growth process as a case of passage or becoming, we can treat it as a unique type of causal process.

One might still disagree with our claim that applying interventionism to CST is appropriate. Suppose that causal set theory is the correct theory of our world. Then it delimits the scope of physical possibility, in this sense: the only physical possibilities that there are, are the ones found in the dynamically acceptable models of the theory. The intervention described in §5.2 involves a dynamically forbidden model (namely,

a finite one) and in this way, the models we are considering are physically impossible. So the intervention we are considering is physically impossible. But, one might continue, interventions ought to be physically possible, and so we have not described a genuine case of causation.

This line of thought trades on a misunderstanding of interventionism. As Woodward makes clear, interventions need not be physically possible (though they should be possible in some broad sense). Here it is worth quoting Woodward at length:

> ... the reference to "possible" interventions in [the interventionist theory] does not mean "physically possible'; instead, an intervention on X with respect to Y will be "possible" as long as it is logically or conceptually possible for a process meeting the conditions for an intervention on X with respect to Y to occur. The sorts of counterfactuals that cannot be legitimately used to elucidate the meaning of causal claims will be those for which we cannot coherently describe what it would be like for the relevant intervention to occur at all or for which there is no conceivable basis for assessing claims about what would happen under such interventions... we thus arrive at the following conclusion: *to a manipulability theory leads unavoidably to the use of counterfactuals concerning what would happen under conditions that may involve violations of physical law.* The reason for this is simply that any plausible version of a manipulability theory must rely on something like the notion of an intervention, and it may be that, for some causal claims, there are no physically possible processes that are sufficiently fine-grained or surgical to qualify as interventions. [33, pp. 132–133]

Thus, assuming Woodward is right, there is nothing wrong with the kind of intervention we have applied to CST, at least from the perspective of interventionism. Granted, it is physically impossible by the lights of that theory. But it is still conceivable, understandable and coherent, and that's sufficient for the interventionist project.

Sixth objection: we have used interventionism to identify actual causation in CST. As noted in §2, however, CST is interpreted to be a non-spatiotemporal theory. But, one might argue, interventionism requires space and time, in this sense: interventions are to be thought of as possible events that happen in space and time that result in some change to a physical system. Moreover, one might argue that we ought to think of interventionism in this way. That's because if we don't, then interventionism is far too liberal, revealing causation where there shouldn't be any.

For instance, one might argue that interventionism can reveal causation even in *mathematical* cases. Take, for instance, the famous seven Bridges of Königsberg. It is impossible to cross all seven bridges, passing over each bridge exactly once. Why? Because if we model the bridges as a graph with each island being a node and each bridge being an edge, then it is possible to prove that the resulting graph will lack an Euler circuit: a path that crosses every edge exactly once. But now suppose we intervene to alter the structure of the bridges. Then there will be such a path. But this intervention is not a causal one: we are considering changes to a mathematical structure, which does not allow for any actual causation to arise. The only way to prevent actual causation from being revealed in this case is to restrict interventions spatiotemporally, thereby preventing their application to pure mathematical structure.

We offer two points in reply to this objection. First, nothing we have said in the paper relies on CST being a non-spatiotemporal theory. Thus, it could be that what the application of interventionism shows is that the theory (with the growth process interpreted physically) is spatiotemporal after all, since it allows for interventions within the dynamics. Perhaps the growth process is a temporal process, and so maybe this is not so surprising. This would not alter our main conclusion, however, that actual causation plays a role in the dynamics so interpreted.

Second, interventionism as it is standardly formulated does not place a condition of spatiotemporality on interventions. This is no part of the theory as Woodward formulates it. Moreover, there is no need to add such a restriction to prevent interventionism from erroneously identifying causation in pure mathematical cases. Instead, we can restrict interventionism in one of two ways. First, we can appeal to the fact that any application of interventionism to pure mathematical cases would involve *impossible* interventions, since mathematical facts are, if true, necessarily true. Thus, we can say that causation only arises when interventions are at least possible (though, as noted, we don't restrict this to physical possibility). Second, we can appeal to the fact that only interventions on physical systems reveal causation, since causation is a physical notion. Thus, the notion of interventionism need not be so liberal as to classify any dependence between variables as causal: only cases that involve possible interventions in physical systems are causal, which are the kinds of interventions we have been considering so far in this paper.

6 Conclusion

We have argued that CST, when situated within an interventionist framework, is causal in a specific sense. In particular, if we assume that the dynamics of CST corresponds to a real physical process, then we can identify relations of actual causation whereby causal set elements bring one another into existence. This shows that, under this interpretation, CST is not just causal in the sense of causation captured by the lightcone structure of general relativity (i.e., causal connectibility). The theory is causal in a deeper sense: there is a sense of causation operative in the dynamics of CST that goes beyond what we find in general relativity and beyond causal connectibility to actual, elemental causation.

As noted in §1, what we have argued is useful. It gives us a better sense of how CST works when we take the growth process physically seriously. In that situation, relations of elemental causal dependence underwrite the dynamics. We've also seen that if we don't take the growth process seriously, then elemental causation may not be playing a role in the theory, though this remains to be seen. Regardless, the analysis we've provided offers a way to interpret the growth process as a causal process.

We also have a sense of the tools available for developing a metaphysical picture of spacetime emergence for this theory. If we interpret the growth process as physical, then we have actual, elemental causation available. This means that the approach to grounding spacetime in causation that we have developed elsewhere is a viable approach to interpreting the theory. It also offers a way of thinking about spacetime emergence as itself a causal process, along the lines of phase transition and other causal cases of emergence.

Finally, on the issue of whether causation plays a role in physics, CST emerges as a potential case study supporting the importance of actual causal relations to physical theory. Whether this case study provides a response to the scepticism about the role of causation in physics is not something we have pursued here, but we note this as a possible option for future consideration.

Acknowledgements The authors would like to thank Sumati Surya for useful discussion of issues surrounding this paper, as well as the audience at the Metaphysics Beyond Spacetime conference in 2024.

Author Contributions All authors contribute to all aspects of the manuscript equally, including but not limited to conceptualisation, design and drafting.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions. Baptiste Le Bihan's work on this article is funded by a Swiss National Science Foundation Starting Grant "Space, Time and Causation in Quantum Gravity", grant number 211317.

Data Availability No datasets were generated or analysed during the current study.

Declarations

Competing interests The authors declare no competing interests.

Ethics approval and consent to participate Not applicable

Consent for publication Not applicable.

References

1. Arageorgis, A.: Spacetime as a causal set: Universe as a growing block? Belgrade Philosophical Annual **29**(33) (2016)
2. Baron, S., Le Bihan, B.: Causal theories of spacetime. Noûs **58**(1), 202–224 (2024)
3. Baron, S., Miller, K., Tallant, J.: Grounding at a distance. Philosophical Stud. **177**, 3373–3390 (2020)
4. Bombelli, L., Lee, J., Meyer, D., Sorkin, R.D.: Space-time as a causal set. Phys. Rev. Lett. **59**(5), 521 (1987)
5. Brightwell, G., Dowker, F., García, R.S., Henson, J., Sorkin, R.D.: Observables in causal set cosmology. Phys. Rev. D **67**, 084031 (2003)
6. Dowker, F.: Causal sets and the deep structure of spacetime. In A. Ashtekar (Ed.), 100 Years of Relativity: Space-Time Structure: Einstein and Beyond, Volume 100, pp. 445–467 (2005)
7. Dowker, F.: Causal sets as discrete spacetime. Contemporary Phys. **47**(1), 1–9 (2006)
8. Dowker, F.: Introduction to causal sets and their phenomenology. General Relativity and Gravitation **45**, 1651–1667 (2013)
9. Dowker, F.: The birth of spacetime atoms as the passage of time. Annals of the New York Academy of Sciences **1326**, 18–25 (2014)

10. Earman, J.: Reassessing the prospects for a growing block model of the universe. Int. Stud. Philosophy Sci. **22**, 135–164 (2008)
11. Frisch, M.: Causal Reasoning in Physics. Cambridge University Press, Cambridge (2014)
12. Hawking, S.W., King, A.R., McCarthy, P.: A new topology for curved space-time which incorporates the causal, differential, and conformal structures. J. Math. Phys. **17**(2), 174–181 (1976)
13. Hitchcock, C.: Probabilistic Causation. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy (Spring 2021 ed.). Metaphysics Research Lab, Stanford University (2021)
14. Hoffmann-Kolss, V.: Three kinds of causal indeterminacy. Australasian J. Philosophy 261–276 (2024)
15. Huggett, N.: Skeptical notes on a physics of passage. Ann. New York Academy Sci. **1326**(1), 9–17 (2014)
16. Huggett, N., Wüthrich, C.: Emergent spacetime and empirical (in)coherence. Stud. History Philosophy Sci. **44**(3), 276–285 (2013)
17. Jaramillo, J.L., Lam, V.: Counterfactuals in the initial value formulation of general relativity. British J. Philosophy Sci. **74**(4), 1111–1128 (2021)
18. Leibniz, G.W.: The metaphysical foundations of mathematics. In L. E. Loemker (Ed.), Philosophical Papers and Letters. Kluwer Academic Publishers. Originally written in 1715 (1989)
19. Major, S., Rideout, D., Surya, S.: Stable homology as an indicator of manifoldlikeness in causal set theory. Classical Quantum Gravity **16**(17), 175008 (2009)
20. Malament, D.B.: The class of continuous timelike curves determines the topology of spacetime. J. Math. Phys. **18**(7), 1399–1404 (1977)
21. Norton, J.D.: Causation as folk science. Philosopher's. Imprint **3**(4), 1–22 (2003)
22. Pearl, J.: Causality: Models, Reasoning and Inference. Cambridge University Press, New York (2000)
23. Price, H., Corry, R.: Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited. Oxford University Press (2007)
24. Reutlinger, A.: Getting rid of interventions. Studies History Philosophy Biological Biomed Sci **43**(4), 787–795 (2012)
25. Rideout, D., Wallden, P.: Emergence of spatial structure from causal sets. J. Physics: Conference Series **174**(1), 012017 (2009)
26. Rideout, D.P., Sorkin, R.D.: Classical sequential growth dynamics for causal sets. Phys. Rev. D **61**(2), 024002 (1999)
27. Russell, B.: On the notion of cause. Proceedings of the Aristotelian Society **13**, 1–26 (1912)
28. Schaffer, J.: Grounding in the image of causation. Philosophical Stud. **173**(1), 49–100 (2016)
29. Sorkin, R.: Geometry from order: Causal sets. Einstein Online 2 (2006)
30. Sverdlov, R., Bombelli, L.: Gravity and matter in causal set theory. Classical Q. Gravity **26**, 075011 (2009)
31. Wilson, A.: Metaphysical causation. Noûs **52**(4), 723–751 (2018)
32. Wilson, A.: Classifying dependencies. In: Glick, D., Darby, G., Marmodoro, A. (eds.) The Foundations of Reality: Fundamentality, Space and Time. Oxford University Press, Oxford (2020)
33. Woodward, J.: Making Things Happen: A Theory of Causal Explanation. Oxford University Press, Oxford (2005)
34. Woodward, J.: Causation in science. In Oxford Handbook of the Philosophy of Science, pp. 163–184. Oxford University Press (2016)
35. Woodward, J.: Causation and Manipulability. In Zalta, E. N., and Nodelman, U. (Eds.), The Stanford Encyclopedia of Philosophy (Summer 2023 ed.). Metaphysics Research Lab, Stanford University (2023)
36. Wüthrich, C.: When the Actual World Is Not Even Possible. In: Glick, D., Darby, G., Marmodoro, A. (eds.) The Foundation of Reality: Fundamentality, Space, and Time. Oxford University Press, Oxford (2020)
37. Wüthrich, C.: The structure of causal sets. J. General Philosophy Sci. **43**, 223–241 (2012)
38. Wüthrich, C., Callender, C.: What becomes of a causal set? British J. Philosophy Sci. **68**(3), 907–925 (2017)
39. Wüthrich, C., Huggett, N.: Chapter 3: The emergence of spacetime from causal sets. In Out of Nowhere. Oxford: Oxford University Press. arXiv:2009.02951 (2020)
40. Wüthrich, C., Huggett, N.: Spacetime from causality: Causal set theory. Oxford University Press, In Out of Nowhere. Oxford (2020)