

Causal Theories of Spacetime

Sam Baron and Baptiste Le Bihan

forthcoming in *Noûs*

Abstract

We develop a new version of the causal theory of spacetime. Whereas traditional versions of the theory seek to identify spatiotemporal relations with causal relations, the version we develop takes causal relations to be the grounds for spatiotemporal relations. Causation is thus distinct from, and more basic than, spacetime. We argue that this non-identity theory, suitably developed, avoids the challenges facing the traditional identity theory.

1. Introduction

Russell once asked:

Can time be derived from causality, or must we retain temporal order as fundamental and distinguish cause and effect as the earlier and later parts in a causal relation? (Russell, 1927, p. 381)

One answer to Russell's question is that temporal relations are based in causal relations. This, in the broadest terms, is the causal theory of time. The causal theory of time has a long tradition of influ-

ential proponents. One of the very first versions of the view was defended by Leibniz.¹ He famously defended a relational theory of temporal succession, according to which temporal succession is reducible to causal relations between material objects. Kant then developed Leibniz's account further, by analysing both relations of succession and relations of simultaneity in causal terms.²

In the twentieth century, the causal theory of time was replaced by the causal theory of spacetime. This theory boasted some prominent advocates, including Reichenbach (1956), van Fraassen (1970) and Grünbaum (1973). Based on pioneering work by Hawking et al. (1976), Malament (1977) and others, it was argued that special and general relativity were, at core, causal theories and the view that the metric structure of spacetime could be accounted for in terms of a causal topology started to gain momentum.³ But the theory was also subject to sustained attack in philosophical circles, especially by Smart (1969), Earman (1972) and Nerlich (1982). The assault on the causal theory was by and large successful, and interest in the theory among philosophers waned.

However, while interest in the causal theory ebbed within philosophy, the core motivations behind the theory never really went away in physics. The work by Malament and Hawking on causal structure in relativity gave birth to an important research programme in physics, culminating in what is now known as *causal set theory*.⁴ As an approach to quantum gravity, causal set theory aims to provide a quantum account of the gravitational field by focusing first on its classical counterpart—as a first step along the way. Core to the causal set theorist's programme is the idea that the metric structure of spacetime can be recovered from its causal structure.

¹ How to classify Leibniz exactly is somewhat open to discussion; see Mehlberg (1980, pp. 48–49).

² See Mehlberg (1980, pp. 54–57).

³ Wüthrich and Huggett (2020a, section 1.3).

⁴ For discussion of the view within physics, see Bombelli et al. (1987); Rideout and Sorkin (2000); Brightwell et al. (2003); Dowker (2013, 2006); Major et al. (2009); Rideout and Wallden (2009). For philosophical work on causal set theory, see Wüthrich (2012); Huggett (2014); Wüthrich and Huggett (2020a); Wüthrich and Callender (2016); Wüthrich and Huggett (2020b).

Given that scientific interest in the causal theory has outlived its philosophical scrutiny, we believe it is time to reignite the approach. The causal theory of spacetime can be broken down into two distinct views, that we call the *identity* and *non-identity* theories. Much of the philosophical criticism of the causal theory has been levelled against the identity theory. However, we will show that the non-identity theory, suitably developed, avoids a core challenge that identity theories face. Note that in developing the causal theory we are assuming that there is at least a minimal notion of dependence within fundamental physics that it makes sense to think of as causation.⁵ In making this assumption, we are working with a broad notion of causation, one that may diverge from causation as understood in certain metaphysical theories. Ultimately, we are not concerned to determine what causation is. If the notion we use here is not considered to be causal by the lights of some approaches to causation, then so be it.⁶ What matters is that the notion we use, and the associated metaphysical picture, can realise the ambitions of philosophers and physicists interested in causal theories of spacetime.⁷

The paper is structured as follows. We begin with a discussion of the identity theory (§2). After that, we introduce a basic version of the non-identity causal theory of spacetime (§3) and develop it so as to avoid the problem facing the identity theory (§4). We complete our development of the view,

⁵ This is controversial. Russell (1927), for instance, argues that there is no causation in fundamental physics, only a notion of ‘functional dependence’ (a view he later repudiated). More recently, Frisch (2014) has argued that there is a viable notion of causation to be found in fundamental physics, one that is closely connected to the interventionist picture we consider in §5.

⁶ There is one version of the objection that we will mention, only to set it aside. The worry is that, intuitively, causation requires time, and so there cannot be fundamental causation without temporal order. We note that the intuition is not universal and has been challenged (see, e.g., Baron and Miller 2015).

⁷ The notion we focus on here is a kind of modal dependence needed to underwrite the version of interventionism that we discuss in §5. This can be thought of as nomic dependence, or simply as a basic modal connection that grounds the rest of the ontology. Whether the picture suggested in this paper is better classified as a causal theory of spacetime or as the heir to the causal programme in the guise of a modal theory of spacetime is a question we leave open at this point.

in §5, by answering two interpretive questions for the theory.

2. The Identity Theory

The identity version of the causal theory can be stated as follows:

The Identity Causal Theory: Spatiotemporal relations between events are identical to causal relations.

The identity theory comes in two forms. The first, and strongest version of the view seeks to define spatiotemporal relations as causal relations, thereby achieving a conceptual analysis of relativistic spacetime in causal terms. The second, more moderate version of the identity theory identifies spatiotemporal relations with causal relations in the manner of the necessary *a posteriori*. Spacetime relations are to causal relations as water is to H₂O.

The strong identity theory faces an immediate difficulty. The trouble comes this way: it is entirely conceivable that there could be a world in which relativistic spacetime exists, but that is completely free of any causal relations. We can, for instance, imagine a spacetime world that is causally idle; in which nothing ever happens. Conversely, we can imagine worlds with causation that do not feature relativistic spacetimes, where space and time are different manifolds (as in Newtonian worlds). These worlds may be physically impossible or even metaphysically impossible, but they are certainly not conceptually incoherent. An *a priori* link between relativistic spacetime and causation, however, would rule out these conceptual possibilities. Assuming that, if *a* can be conceived of without *b*, then the identity of *a* and *b* is not *a priori*, there does not seem to be a conceptual link between causation and relativistic spacetime of the right kind.

A second difficulty for the strong identity theory is voiced by Smart who complains that:

To elucidate the concept of space-time in terms of the concept of causal connectedness seems to be to elucidate the comparatively clear by reference to the comparatively unclear.

(Smart, 1969, p. 394)

It is tempting to dismiss this objection as a product of its time. Smart made this comment before much of the contemporary work on the metaphysics of causation had been written. Moreover, verificationist suspicions about metaphysics in general (and causation in particular) were still thick in the air. Given that our understanding of causation has progressed a great deal since then, Smart's worry has less bite. But we do not think it is completely toothless. For it is still correct that we have a very mathematically precise understanding of relativistic spacetime in geometric terms. While we have developed theories of causation that have some level of precision, what we take to be the most precise of these—the interventionist account coupled to the structural equation framework—typically foregoes any reductive ambitions and takes causation to be an unanalysed primitive. Arguably, causation is still less well understood than spacetime.

The weak identity theory avoids the broadly conceptual problems voiced above, and has some claim to being the standard form of the view within philosophy. The weak identity theory does not require one to define spatiotemporal relations in terms of causal ones, any more than the identification of water with H₂O requires one to strictly define water in molecular terms. Just as the concepts of water and H₂O can differ radically, so too for the concepts of spacetime and causation. Since one is not seeking any kind of conceptual analysis, it also doesn't matter whether causal relations are less well understood than spatiotemporal ones, the identification might still be apt (and useful).

But the weak identity theory faces a serious problem of its own. A version of the difficulty is (again) raised by Smart. The problem, as Smart puts it, concerns the possibility of spatiotemporally located entities that are neither causes nor effects. Smart writes:

It is difficult to see how the causal theory of time is applicable to theories which allow for the existence of events which are neither causes nor effects of other events. It at least seems to me that I can consistently envisage a universe of purely random events spread out through space-time. (Smart, 1969, p. 394)

According to the weak identity theory, spatiotemporal relations and causal relations are identified. Thus, any spatiotemporal relation must be a causal connection, which means that there cannot be any entities that are causally idle so long as they bear spatiotemporal relations to other entities (which they must if they are located in spacetime). The weak identity theory implies, incorrectly according to Smart, that everything does and must do something.

Call this the problem of causal indolence. The problem is really a cluster of three worries that are not usually distinguished. The three worries are based on three different situations: (i) empty spacetime regions; (ii) timelike, but not causally, connected events and (iii) spacelike, but not causally, connected events. Let us briefly consider each in turn, in order to get a better sense of the difficulty.

Regarding empty spacetime regions, the worry is that there might be regions of spacetime that are completely free of matter and energy. Within such regions, there doesn't seem to be anything to connect via causal relations. Accordingly, there do not seem to be any causal connections to which the relevant spacetime relations might be identified.

This version of the worry should not bother us for at least two reasons. First, the notion of an empty spacetime region suffers from a crucial ambiguity. On the one hand, an empty region of spacetime might be a region of spacetime that lacks any matter fields at all. On the other hand, an empty region of spacetime might be a region of spacetime throughout which matter fields are distributed, but the field values for the matter fields are all zero for that region.

It is plausible that there are empty regions of spacetime in the second sense (at least on average, neglecting quantum fluctuations). That there are empty regions of spacetime in the first sense seems simply wrong in the context of general relativity. In general relativity, wherever there is a metric field, there are matter fields there as well. It may just be that the matter field values are zero (and thus the spacetime is locally flat and, seemingly, empty of interesting material content).

Being ontologically serious about fields means accepting that there is always something, even when the value of the field is zero. More precisely, there are determinable and quantitative properties

that can act as relata for causal relations. Accordingly, we have more in our ontology than a cursory glance might suggest. It is now possible to see how this ontological basis could be interpreted as causal. As we will explain below, the interventionist framework can be used to detect the presence of causal relations in terms of dependence relations between events. Within this framework, parts of a matter field can be viewed as handles for interventions even within a vacuum solution to the field equations, in which all matter fields are zero-valued everywhere. For even in this case, if the value of the field were to change from zero to another value at a particular location, subsequent values of the field at nearby timelike distances in the future would also be affected. Since, within the interventionist framework, the presence of interventions of this kind can be used to establish the presence of causal relations, there is scope to treat the matter field as causal even when it is zero-valued.⁸

The second reason to resist the empty spacetime worry is that it is unclear that the matter fields can have a zero value anywhere in the actual world. The remnant light of the Big Bang—the so-called cosmic microwave background radiation—permeates the universe everywhere. This radiation is an oscillation of the electromagnetic field—a matter field—and so at least this field is not zero-valued anywhere. So there really doesn't seem to be a sense in which spacetime is 'empty', at least not in a manner that would undermine causation.

This brings us to the second kind of situation: situations wherein timelike separated events are not further related by causal relations. The idea is *prima facie* quite simple. Many pairs of events are timelike separated—one can be represented as located in the past or future light-cone of the other. Among those pairs, only some of them are actually causally connected. Think for instance about you raising your hand right now and the burst of a supernova in galaxy Andromeda in the far distant

⁸ As we will explain below, one may not be willing to treat interventionism as a basis for causation at the fundamental level. Even if one holds such a view, there are still modal connections of the kind that support interventions, and so there will be a way to develop a 'modal' theory of spacetime that, while not causal by the lights of some philosophers, is still within the spirit of the causal theory of spacetime, broadly construed.

future: those two events do not seem to be causally connected in the ordinary sense of causation. However, they are timelike separated.

As with empty regions of spacetime, it is unclear that situations in which there are timelike, but not causally, separated events can genuinely arise. As previously discussed, there is a sense in which matter fields are everywhere (such as the electromagnetic radiation from the Big Bang). If that is the case, then actual signals are systematically propagating between any timelike separated events (namely between parts of the electromagnetic field as it oscillates). Those relations can be trivial and make not much of a difference. Still, the differences those relations make, however tiny, are causal in nature, and that is all that is really needed for the weak identity theory. Again, think of you raising your hand right now and the burst of a supernova in galaxy Andromeda in the far distant future; there is a sense in which the event in Andromeda will be affected in a negligible way by the modifications of the field caused by you raising your hand.

The third and final situation is the most worrying: the relations between spacelike separated events. Those events, in order to be causally related, would need signals propagating faster than the speed of light in a vacuum, which is physically impossible according to the general theory of relativity. Hence, those events are spatiotemporally related but they cannot, as a matter of principle, be related by a physical signal of any kind. This is widely held to show that there are no causal relations between such events.⁹ If true, it follows that there are no such relations with which the relevant spacelike separations can be identified.

The problem of causal indolence, then, is strongest when framed as a worry about spacelike separation. Moreover, this version of the worry is untouched by the considerations adduced thus far

⁹ It could be that some spacelike related events are sometimes causally connected, as suggested by some approaches to quantum mechanics. We set aside the possible existence of such spacelike connections; arguably, even though *some* spacelike separated events could be causally connected, it will not be true of *all* spacelike separated events. Thus, the generic problem of spacelike but not causally related events would not be addressed.

for the other versions of the worry. Even if matter fields propagate throughout the entire universe, still there cannot be any causal connections between spacelike separated components of the field.

Can the problem of causal indolence be solved? The standard strategy for dealing with the problem is to trade actual causal relations in for a notion of causal connectibility (see Grünbaum 1973, p. 187). Causal connectibility is a modal relation. Roughly, x is causally connectible to y when it is possible for there to be a causal relation between them.

At best, however, the appeal to causal connectibility can help us to address the first two versions of the causal indolence problem. But we have already suggested that these versions of the problem are not that pressing. The real trouble lies with spacelike separation. Here, causal connectibility is no use as spacelike separated events are not even causally connectible in an important sense. Indeed, in no solutions of general relativity are there spacelike separated events related by physical signals. Spacelike separations, in this context, are generally viewed as excluding causal connections.

One could claim that such events are connectible: they could be related if some nomological constraints were lifted, or, if causation at a distance were admitted in one's ontology. But it is not clear what a notion of connectibility that goes beyond physical law amounts to and the cost of causation at a distance is not negligible. Accordingly, we recommend looking for another way to address the problem of causal indolence, a solution that does not appeal to a primitive notion of causal connectibility between spacelike separated events.

3. Causal Set Theory

The causal indolence problem motivates a shift away from a version of the causal theory of spacetime based on identity. In order to get a sense of how we should reformulate the causal theory of spacetime, it is instructive to take a closer look at causal set theory. As briefly noted in §1, causal set theory is a programme in physics that continues to uphold the core motivations that led to the causal theory of spacetime in the first place. Causal set theory is important in the current context because it inspires a

new way of thinking about the causal theory of spacetime.

At the most fundamental level, physical reality in causal set theory is constituted by causal sets (for an overview, see Dowker 2006). Each causet can be represented as a pair $\langle E, \leq \rangle$ in which E is a set of elements and \leq is a relation of causal precedence defined on E . Causal set theory obeys two axioms (Wüthrich 2012, p. 229). First, \leq is defined as a partial order. Second, causal sets are locally finite (which captures their discreteness). If we let $|E|$ be the cardinality of E , and if we let \aleph_0 be the cardinality of the integers, then the second axiom can be expressed as follows:

$$\forall xyz \in E, |\{x \leq y \leq z\}| < \aleph_0$$

Causal sets form the kinematical core of causal set theory. These sets are generally supplemented with a law of sequential growth (Rideout and Sorkin 2000), which imbues causal set theory with a dynamics. Roughly speaking, the law of sequential growth corresponds to a ‘process’ whereby causets ‘grow’ as new elements ‘come into existence’ and are connected to existing elements via \leq .¹⁰

We have put ‘grow’ and ‘come into existence’ in scare-quotes, because these words have a temporal if not spatiotemporal flavour to them. It is not at all clear, however, that the dynamics of causal set theory should be interpreted in temporal or spatiotemporal terms. There is no time dimension in which causal set elements ‘come into existence’ one after another. If there were such a time dimension, then it should be the case that the causal set elements stand in temporal ordering and, perhaps, metric relations to one another. However, the fundamental structure described by causal set theory does not seem to be spatiotemporal. As Huggett and Wüthrich put the point:

There simply is nothing on the fundamental level corresponding to lengths and durations (or, more generally, to spacetime intervals), and no alternative interpretation of the causal sets in terms of metrical relations is available. (Huggett and Wüthrich, 2013, p. 278)

¹⁰ The process is a Markov process: a sequence (such as a sequence of causal set elements) in which the probability of each member within the sequence depends only on the probability of the member before it.

At this point, it is worth pausing to consider an objection. We have said that causal set theory features a kinematics and even a dynamics. However, it also seems to lack spatiotemporal structure. One might worry, however, that kinematics and dynamics presuppose spatiotemporal notions, and thus that these aspects of the theory cannot really be understood in the absence of spacetime. We admit that we are adopting a particular interpretation of causal set theory, according to which kinematics and dynamics can be understood without spacetime. Exactly how to understand kinematics or dynamics without spacetime is highly dependent on how the physics of causal set theory develops. We thus leave it as a question for physicists as to how we might make sense of the kinematical and dynamical aspects of a theory without spacetime. Our project here is to simply proceed on the assumption that kinematical and dynamical notions can be freed from their usual spatiotemporal understanding, and then see how we might understand the metaphysics of such a view.

At any rate, despite the lack of spatiotemporal relations at the fundamental level, it is nonetheless expected that spacetime will be an emergent phenomenon in causal set theory. It is hoped that certain sufficiently complex arrangements of causal sets will correspond to spacetime structure once we ‘zoom out’ from the scale of individual causal set elements.

There are some preliminary results that suggest spacetime will, indeed, be emergent in the context of this theory. For it has been shown that there are certain causal set structures that can be mapped into the metric structure of relativistic spacetime (see Wüthrich 2012 and Dowker 2013 for an overview). The key is to use a specific probability function that ‘sprinkles’ causal sets so that, on average, the number of causal sets in a given portion of the total structure defined by the partial order corresponds

to the volume of a region of spacetime.¹¹ What the sprinkling procedure produces is a causal set $\langle E, \leq \rangle$ such that there is a map $\phi, \langle E, \leq \rangle \rightarrow \langle M, g \rangle$, where M is a manifold, g is a spatiotemporal metric, and ϕ satisfies the following three conditions (we use the formulation from Wüthrich 2012, p. 232):

1. The causal relations are preserved, as follows: $\forall a, b \in E, a \leq b$ iff $\phi(a) \in J^-(\phi(b))$, where $J^-(p)$ is the causal past of p , i.e. the set of points $q \in M$ such that there exists a future-directed causal curve from q to p (or $q = p$).
2. On average, ϕ maps one element of E onto each Planck-sized volume of $\langle M, g \rangle$.
3. $\langle M, g \rangle$ has no length scales smaller than the discreteness scale of the causal set; for instance, it is approximately flat below the Planck scale (the typical discreteness scale).

Note that not all causal sets can be mapped to spatiotemporal manifolds. Indeed, causal set theory, at least at the mathematical level, allows many causal sets not to have any manifold-like approximation. Perhaps some of these causets can be ignored as mathematical pathologies. However, not all of those non-spatiotemporal causal sets are expected to be artefacts of the mathematical structure of the theory, to be eventually ruled out. ‘Physical’ but not spacetime-like causal sets might well be

¹¹ What we need is a uniform distribution of the elements of the causal set across the manifold. The sprinkling process is a particular technique that ensures the preservation of Lorentz invariance in the relativistic description, when distributing the elements. Note that the CST structure itself, being discrete, cannot be Lorentz invariant; but the corresponding relativistic spacetime must satisfy this constraint. Such a distribution cannot be obtained by ‘cutting’ the manifold into minimal four-dimensional volumes and assigning to each minimal volume an element of the causal set, as this would already compromise Lorentz invariance in the relativistic description. Indeed, any cutting of this sort will require the specification of a preferred length scale. The sprinkling procedure being random, ensures that the way the discrete structure is built does not violate Lorentz invariance from the start. Finally, it should be noted that the sprinkling procedure may not be the only way to get around this issue. But it is the actual way physicists working on CST managed to bypass the difficulty.

needed to describe the high-energy quantum domain (see e.g. Zalel 2020, p. 2). If so, then the causal theory suggested by the causal set approach can hardly rely on a straightforward identification of spatiotemporal relations with the causal relations that are constitutive of causal sets.¹²

Actually, this last point may not seem entirely obvious. For consider again the first condition of the sprinkling procedure described above that shows how to produce spacetime structure from a causal set:

1. The causal relations are preserved, as follows: $\forall a, b \in E, a \leq b$ iff $\phi(a) \in J^-(\phi(b))$, where $J^-(p)$ is the causal past of p , i.e. the set of points $q \in M$ such that there exists a future-directed causal curve from q to p (or $q = p$).

Doesn't this just tell us that the causal relations between causal set elements must be placed into a correspondence with the causal relations in spacetime? Not quite. The notion of the 'causal past' and 'causal future' specified here is none other than the backwards and forwards light-cones of a spacetime event. We could thus remove any mention of the causal future and past and replace it simply with the set of timelike connections going into and out of a given event. That is, we can rewrite the above condition by replacing $J^-(\phi(b))$ with the timelike past of an event $I^-(\phi(b))$. The procedure would work just the same, since what matters is whether we can produce the timelike structure of spacetime. It just so happens that, in this procedure, causal history is being used as a proxy for timelike structure. Ultimately, though, that timelike structure is still an emergent aspect of the theory, and not a fundamental connection between causal sets.

One important question is whether ϕ is unique in the following sense: is there some set of conditions on ϕ such that for any mapping ψ that satisfies those conditions, if ϕ maps a causal set $\langle E, \leq \rangle$

¹² Furthermore, other approaches to quantum gravity suggest that spacetime will be non-fundamental in a similar way, making it even more pressing to come up with a causal theory of spacetime that does not require the strict identity of causation and spacetime. For a review, see Huggett and Wüthrich (2013) and Crowther (2018).

to a manifold pair $\langle M, g \rangle$, then ψ maps $\langle E, \leq \rangle$ to a manifold pair $\langle M^*, g \rangle$ such that $\langle M, g \rangle$ and $\langle M^*, g \rangle$ agree in relevant spatiotemporal respects? Or, to put the point slightly differently, does ψ map each causal set structure $\langle E, \leq \rangle$ to a single manifold pair $\langle M, g \rangle$, or can the same $\langle E, \leq \rangle$ be mapped to multiple manifold pairs, each of which disagree in important respects?

If ϕ is unique in the relevant sense, then it is tempting to say that spatiotemporal structure is, in some good sense, a redescription of a causal set structure. This would thus tend to lead us back toward something like an identity version of the causal theory of spacetime, albeit perhaps not exactly the kind of identity theory that has been articulated thus far.

Ultimately, we are open to the possibility that the mathematical relationship between the causal set structure and spacetime may be such that it makes sense to view spacetime as a redescription of the causal set structure. In this case, we would perhaps recommend resuscitating an identity theory. However, as matters stand, it seems that ϕ is not unique, in the relevant sense. While it is conjectured that the same causal set can be mapped onto spacetimes that are similar, these similar spacetimes can diverge in spatiotemporal respects. These divergent spacetimes will ultimately look highly similar at a certain scale, but not below that scale.¹³ But what this scale is exactly, and how to define it, is difficult, and so the conjecture is not even well defined. Interestingly for our purpose, even if the conjecture were to hold, there would still be no straightforward path from CST to the identity theory, as the mapping between CST solutions and GR models via ϕ won't be one-to-one—a condition we

¹³ Note that the scale below which spacetimes are similar is not the Planck scale, but rather a scale between the Planck scale and some other, broader, scale.

take as a necessary requirement for the identity theory.¹⁴ Again, though, whether there is a stronger mathematical relationship between causal sets and spacetime remains open.

The emergence of spacetime within causal set theory suggests a very natural way to reformulate the causal theory of spacetime. Emergence can sometimes signal a kind of ontological dependence, whereby the emergent entity depends on something that is more fundamental. So, for instance, water is an emergent entity at macro scales. This is partly analysed in terms of a dependence relation between water and more fundamental H₂O molecules: water is grounded in certain molecular structures. Similarly, rather than taking spatiotemporal relations to be identical to causal relations, as per the identity theory, we can take spatiotemporal relations to be grounded in, though ultimately distinct from, causal relations. We can state this version of the causal theory as follows:¹⁵

The Non-Identity Causal Theory: Spatiotemporal relations between events are grounded in causal relations.

The notion of grounding in the non-identity theory is left open, except insofar as we take grounding to imply a failure of identity. We use the term ‘grounding’ as a general way of referring to a variety of possible relations that may or may not align with the standard use of the term—we have in

¹⁴ We take it that an identity theory would require that for a causal set structure $\langle E, \leq \rangle$, and for any manifold pairs $\langle M, g \rangle$ and $\langle M^*, g \rangle$ such that $\langle E, \leq \rangle$ is mapped to $\langle M, g \rangle$ and $\langle M^*, g \rangle$ via ϕ , then $\langle M, g \rangle$ and $\langle M^*, g \rangle$ will exhibit the same metric at all spatiotemporal scales (which we take to be all physical distance scales above the discreteness cut-off). This is precisely what fails even if the conjecture described above is true. $\langle M, g \rangle$ and $\langle M^*, g \rangle$ will agree at some scale S which is not yet defined, but for scales between the cut-off point and S , $\langle M, g \rangle$ and $\langle M^*, g \rangle$ will not agree in all metrical respects.

¹⁵ A note on terminology: we continue to use the word ‘event’ to specify the relata of causal relations. We recognise that the word carries some spatiotemporal connotations and so may not be entirely appropriate once we have shifted to a non-identity theory, but it makes the discussion neater. Strictly speaking, we define an event as any element within a causal structure, in the sense discussed in §4.

mind a very broad understanding of grounding that includes any ontological dependence relation that could relate two sets of numerically distinct entities. Different ways of understanding grounding will give rise to different versions of the non-identity theory. What all of these views have in common, however, is that causal relations are taken to be more fundamental than spatiotemporal ones.

Does a shift to the non-identity theory help us to overcome the causal indolence problem? Recall that the crux of the causal indolence problem concerns spacelike separated events. The difficulty is that for these events, there are no causal relations (and no merely possible causal relations) that connect them.

In the case of the non-identity theory, we are not required to find causal relations between spacelike separated events that can be identified with the spatiotemporal relations between them. So the non-identity theory boasts at least a *prima facie* advantage over the identity theory when it comes to causal indolence.

That being said, we are still apparently tasked with finding causal relations that can *ground* spacelike separated events. On a naive reading of the non-identity theory, however, the prospects for doing that may seem equally bad. For when there are no causal relations between spacelike separated events, it is difficult to see how causal relations could serve as the ontological basis for the spatiotemporal relations at issue (for the same reason that it is hard to find causal relations to identify spacelike connections with). In order for the non-identity theory to be viable, then, more must be said concerning the manner in which spatiotemporal relations are grounded.

4. The Non-Identity Theory Refined

Our goal in this section is to provide a more refined version of the non-identity theory. The core of our approach is informed by the observation that grounding is a more flexible relation than identity. Unlike identity, grounding allows for a relation at one level to be grounded in the absence of a relation at another level. To take a very simple example, suppose that there is a private club with the following

membership conditions: someone can only be a member of the club if they are not biologically related to anyone already in the club. In this situation, the membership relation between two members of the club holds, at least in part, in virtue of the absence of any underlying biological relationship between them.

Another way in which grounding is more flexible than identity concerns the ‘size’ (for want of a better term) of the grounds compared to the grounded. Grounding allows that what underlies a relation might be more than just a single relation. What grounds a relation might be a great many relations, or a complex system of facts. Take, for instance, the relation of citizenship that many people stand in toward a country. This relation holds, when it does, because of an entire complex of facts concerning the individual and the country at issue. There is no single factor that grounds citizenship.

Our proposal is to use these two dimensions of grounding to develop the non-identity theory. To a first approximation, our view is that spatiotemporal relations are grounded in a pattern of more fundamental causal relations between events. Roughly, then, spatiotemporal relations between timelike separated events are grounded in the *presence* of causal relations between events. Thus, if two physical events are linked by a fundamental causal relation, then that is sufficient for them to be timelike connected as well. Relations between spacelike separated events, by contrast, are grounded in the *absence* of causal relations between events. In particular, if two physical events are not linked by a fundamental causal relation, then that grounds a spacelike connection at the spatiotemporal level.

This version of the non-identity theory needs to be developed in two ways. First, the absence of causal connections promises to over-generate spacelike connections. Suppose, for instance, that there are mathematical objects. Mathematical objects are generally taken to be causally inert: standing in no causal connections to physical events. If the mere absence of causation is sufficient for (some) spatiotemporal connections, then mathematical objects (if they exist) will be spacelike connected to physical events. Of course, the case of mathematical objects just dramatises the problem: any two things that aren’t causally connected will be spacelike connected unless further constraints are placed

on the grounding base.

Second, one might worry that the absence of causal connections, on its own, is too modally weak. The mere absence of a causal connection is compatible with its possible presence. As discussed, however, spacelike separated events are ones that cannot be causally connected via a physical signal. We should expect this to be represented in the fundamental causal structure of the grounding base for spatiotemporal relations.

We can address both of these issues by expanding the grounding base for spatiotemporal relations. To begin with, we can say that in order for x and y to be spatiotemporally connected at all, they must both be parts of a causal structure. A causal structure is a set of elements in which every element is connected to at least one other element by some causal connection (we continue to use the term ‘event’ for elements embedded in a causal structure in this manner). Our picture forbids *isolated events*, events that bear no fundamental causal connections to *anything*. A causal set is an example of a causal structure of this kind: there are no elements in a causet that are not ordered with respect to some element by \leq .

By expanding the grounding base for spatiotemporal relations between events to include facts about being part of a causal structure, there is no fear that mathematical objects will end up being spatiotemporally connected to physical events. That’s because mathematical objects are not parts of any causal structures, because they do no causal work. Spacelike separated events, by contrast, may not be causally connected to each other, but they are causally connected to other events, and so are parts of a causal structure.

If we make the further assumption that causal structures are rule governed, then the issue concerning modal strength can be addressed as well. Causal structures are rule governed just when there are physical laws that dictate not just which events happen to be causally connected within that structure, but also which events *can* or *cannot* be connected. Thus, a causal structure in the relevant sense is not just a pattern of actual causal connections, it is also a pattern of causal connectivity.

Once we have added a modal dimension to the grounding base we can say quite generally that whether two events are spatiotemporally related, and how, depends on the presence or absence of causal connections between those events, along with the modal strength of those presences and absences. This gives us a greater flexibility to handle both spacelike and timelike connections between events that may not be causally connected at a fundamental level. Thus, we can say that when two events are timelike separated that is always just because there is a fundamental causal connection between them.¹⁶ When two events are spacelike separated, by contrast, that is because there is no fundamental causal connection between them and there cannot be any such causal connections, given the overarching rules of the causal structure in which those events are embedded. We assume that the laws governing causal structures are more fundamental than the laws captured by general relativity, whatever those laws might be. The pattern of modal relationships within a causal structure, then, is given by these more fundamental laws. In saying this, we are not committing ourselves to the view that the laws of general relativity are not really laws, because they are not fundamental. We are willing to be quite permissive about what does or does not count as a law since nothing for us hangs on this.

How are we to ground the metric connections in the ontology of causal set theory? Note that we are aiming to ground physical spatiotemporal distance relations represented in a more fundamental structure. In particular, what we are aiming to provide is a grounding story of the following broad kind: for spacetime points x , y and z , related via a spatiotemporal distance interval d , if $d(x, z) > d(x, y)$, then there are fundamental elements in the causal set structure (or groups of elements) a , b , c such that $d(a, b) > d(a, c)$. Thus, it is not the precise values for distance between points that we want to ground, but rather the relative distances between points.

¹⁶ The distinction between timelike and null connections can be handled in the usual way. Two events are null connected when the only possible causal relations between those events are relations involving massless particles such as the emission of light from a source causing the absorption of light at a detector.

Note also that we are not aiming to ground *every* point that we find in the mathematical description of the metric field of general relativity. The metric field in general relativity is continuous, and so there are many more points in the metric field than there are causal set elements. It will thus be difficult to find enough elements in the causal set structure to ground the relative distances between all points in the manifold description. Rather, what we aim to do is ground spacetime points up to a certain scale factor, where that scale factor is a proposed discreteness cut-off (such as the Planck scale). What this means is that we are assuming—as seems natural in quantum gravity—that physical spacetime relations breakdown at a certain scale, and thus that there is surplus mathematical structure in the description of the metric field of general relativity that does not correspond to any physical spatiotemporal distance relation. Despite this, we still take spacetime points at a certain scale to correspond to causal set elements (or groups thereof). It is just that below this scale, there is no such correspondence to be found. The correspondence itself, however, is an important part of how we recover metric connections, as we discuss below.

Finally, note that in our grounding story we aim to preserve the split between spacelike and timelike relations. Accordingly, we ground spacelike relations and timelike relations differently, at the fundamental level, in different aspects of the causal set structure. We thus have two quite different cases: timelike and spacelike directions. As the grounding of timelike relations is more straightforward than the grounding of spacelike relations, let's start with timelike relations. In outlining the grounding story, we will talk of correspondence between spacetime points and elements of the causal set. As we discuss below, this is a simplification of the picture, one that we adopt for the ease of exposition.

Timelike distances between two spacetime points are grounded in the number of causal links between two causal set elements. The distance between two spacetime points in the relativistic description thus corresponds to the shortest path through the causal network, between two causal set elements corresponding roughly to the two spacetime points.

Spacelike distances do not correspond to any direct connection between the causal set elements that correspond to those points. However, the causal set elements that ground spacelike distances are indirectly connected. This indirect connection between the two causal set elements goes through a common ancestor. The distance between any two spacelike separated points is then grounded in the minimum number of relations of causal precedence it takes to arrive at a common causal ancestor within the underlying causal set description. In other words, the spacelike distance in the relativistic description is grounded in another notion of distance, expressed by the number of relations of causal precedence structuring the shortest path back to a common causal ancestor between two points.

So, for instance, for spacelike separated points x and y that correspond to causal set elements a and b , the spacelike distance $d(x, y)$ corresponds to the shortest path one can take through the causal set structure to arrive at a common ancestor of a and b , c . This allows for the grounding of relative spacelike distance as follows. For spacetime points x , y and z that correspond to causal set elements a , b and c respectively, $d(x, y) > d(x, z)$ when the shortest distance it takes to arrive at a common ancestor of a and b in the causal set structure is greater than the shortest distance it takes to arrive at a common ancestor of a and c in the causal set structure. A case along these lines is depicted in Figure 1.

Note that since the causal set structure is a partial order, there will always be a common ancestor in the structure for any two elements (this is just a reflection of the point noted above that there are no elements in a causal set structure that fail to be connected to at least one other element). Note also that this example is simplified in a certain respect: it assumes that each spacetime point can be mapped to a single, corresponding element of the causal set structure. In general, we see no reason to suppose that the correspondence will be this simple. Rather, the correspondence may be more complex, either involving a correspondence between regions of spacetime and causal set elements, or between regions of spacetime and collections of causal set elements. This added complexity doesn't change the underlying story of how metric connections are grounded, however, so long as relative

distances between spacetime points or regions correspond to relative distances between causal set elements or groups of elements.




Figure 1: *Grounding Metric Connections.* On the left, the number of CST elements ground the timelike relations in relativistic models. On the right, the shortest path between two CST elements and their common past ground the spacelike distance in relativistic models.

The left and right images in Figure 1 depict the grounds for spacelike and timelike distances respectively. On the right, the distance between two points represented by the triangle is traced back to a common ancestor, also represented by a triangle. This distance is smaller than the distance between points represented by the square, which is also traced back to a common ancestor, also represented by a square. The difference between these distances grounds the difference between two spacelike distances, assuming a mapping from parts of spacetime to causal set elements. On the left, the distance between two points depends on how many links in the causal structure are between them. This difference in the number of links grounds differences between timelike distances between points, again assuming a mapping from parts of spacetime to causal set elements. In both cases, the smallest distance is used: spacelike distances are grounded in the smallest distance back to a common ancestor in the causal set structure, and timelike distances are grounded in the smallest

number of links between points in the causal set structure.¹⁷

One further refinement of the non-identity theory is in order. Recall that, in causal set theory, not every causet gives rise to an emergent spacetime. Only some do. Thus, it must be that the causal structures that ground spatiotemporal relations must be of a particular type in order to perform the relevant grounding role.

Exactly what type or types of causal structures give rise to spatiotemporal relations is not a question we can hope to answer here. For, presumably, which types of causal structure give rise to spacetime depends heavily on the relationship between the laws that govern the causal structure at issue, and the non-fundamental laws of general relativity. Determining which types of causal structures give us spatiotemporal relations must therefore await further developments in physics. For now, however, we can capture this aspect of spacetime emergence by specifying a general restriction by type for causal structures. Only when the structure is of the right type, where that means that it is governed by the right laws, will spacetime emerge.

We can now summarise the grounding story for spatiotemporal relations as follows:

The Refined Non-Identity Causal Theory: Spatiotemporal relations are grounded in the presence or absence of fundamental causal connections between events that are embedded in a causal structure C where (i) C is governed by laws that qualify it to be a grounding base for spacetime and (ii) the laws on C determine the modal status of each presence or absence of a causal connection within that structure.

Before proceeding it is worth considering an objection. We have said that spatiotemporal relations are partly grounded in *absences*. For example, the presence of a spacelike relation at the spacetime

¹⁷ Note that only certain causal set structures will have the right shape to ground spacetime. We presume that the the causal set structures that are capable of grounding spacetime will be the very same as those delivered by sprinkling causal set elements into a manifold via a Poisson process.

level is partly grounded in the absence of any causal connections at the more fundamental level of a causal structure. One might worry, however, that there cannot be any grounding by absence. Grounding, one might argue, must always obtain between the presence of an object, and the presence of a more fundamental object.

Note, however, that the grounding of spatiotemporal relations is never entirely a matter of absence. Even spacelike connections are partly grounded in the presence of events that are embedded in a causal structure. We see little reason to be concerned about this kind of grounding. As previously discussed, there are clear examples of grounding by at least partial absence in this sense (as in the club membership example).

If one remains worried about grounding by absence, then it will be difficult to find a grounding base for spacelike connections. As we see it, there are two options. One option is to assume that, at the more fundamental level of a causal structure, all events are causally connected, even events that are spacelike separated. We can then draw a distinction between two kinds of causation: fundamental causation, which underwrites spacetime, and non-fundamental causation, which arises within spacetime, and thus which exists at the spatiotemporal level. This distinction allows us to say that while all events in spacetime—even spacelike separated ones—are causally related to one another at the fundamental level, only some of those fundamental causal relations ground non-fundamental causal relations. In particular, while fundamental causal connections ground causal relations between timelike separated events at the spatiotemporal, non-fundamental level, they do not ground causal connections between spacelike separated events at that level.

The relationship between spatiotemporal causation and fundamental causation may be deemed roughly analogous to the relationship between causation at the macro level and causation at the micro level. Not every micro-level causal connection issues in a macro-level causal connection, since micro-causation can get ‘washed out’ at the macro level. Moreover, at the macro level there can be entirely new causal relations that, while grounded in micro-causal detail, are nonetheless unique to the macro

level. The divide between spatiotemporal causation and fundamental causation is a deeper, more enduring split than the macro/micro split, because there is spacetime at one level and not at the other. But the same basic picture applies: spatiotemporal causal relations display a degree of autonomy from fundamental causal relations, even while being grounded in them. This is a familiar feature of causal emergence.

The distinction between fundamental and emergent causation allows us to capture the sense in which causal connection between spacelike separated events is forbidden by relativity, since there will be no such causal connections within spacetime. On this picture, however, the distinction between spacelike and timelike relations is not metaphysically deep. It is an artefact of emergent causation, and does not show up at the fundamental level.

If we reject both grounding by absence and the notion of emergent causation, then the only remaining way to avoid the problem of causal indolence would be to combine the non-identity theory with anti-realism about spacelike connections. On this account, spacelike connections simply do not exist. The relativistic manifold representation adds mathematical surplus to the physical spacetime which does not, generically, include spacelike relations.

An anti-realist approach to spacelike connections is available to the proponent of an identity theory as well. So modifying the non-identity theory in this way would weaken the motivations for moving away from the identity theory in the first place. We would still have reason to give up the identity theory, however. That's because, as noted in the previous section, the selective emergence of spacetime from causal structures in causal set theory is difficult to square with the identification of spatiotemporal relations with causal ones. Ultimately, however, we wish here to remain realist about relativistic spacetime taken as a whole in the context of the grounding approach. While we have a certain affinity for anti-realism about spacelike relations, we believe that this view should be further explored in the context of a discussion of the weak identity theory, and set it aside in the context of the discussion of the grounding view. We set aside the emergent causation option as well since we

don't see any reason to be worried about grounding by absence.

5. Two Interpretive Questions

So far we have outlined an alternative to a causal theory of spacetime based on identity. The alternative appeals to grounding as the central relation connecting spacetime to a more fundamental causal structure. As it stands, there are two interpretive questions that remain open for the theory. First, what is the grounding relation supposed to be? Second, what does it mean to say that causation is fundamental? Different versions of the non-identity theory can be produced depending on how these questions are answered. For the most part, we want to leave these questions open, so that a wide-range of causal theories of spacetime can be developed using our approach as a basis. That being said, it is perhaps instructive to see how these questions might be answered, and thus how a non-identity theory of spacetime might be fleshed out. Thus, in what remains we offer our preferred answers to both questions.

5.1. Grounding and Mereology

We have said that spatiotemporal relations are grounded in more fundamental causal relations between events. As discussed, we take grounding to be a generic way of referring to some more specific relation. The more specific relations correspond to what Bennett (2017) refers to as 'building relations', which include (at a minimum) mereological composition and material constitution.

Our preferred understanding of grounding in this context is mereological. In this way, the refined non-identity theory can be turned into a class of mereological models. In what follows we sketch out a class of simple models along these lines. Our simple models draw on the notion of *logical* mereology developed by Paul (2002, 2012). In rough terms, logical mereology is the extension of mereological concepts to properties and relations. Thus, just as we might say that a particular object

has parts, so too can we say that properties and relations have parts. This allows us to construe the grounding relations between spatiotemporal relations and more fundamental causal connections in mereological terms.

In order to build a simple mereological model of the non-identity theory, we start with a quintuple $\langle E, R, C, P, F \rangle$. E is a set of events. R is a set of spatiotemporal relations between events that can be partitioned into timelike and spacelike relations. C is a primitive causal relation, and is a partial order on E . P is a primitive parthood relation that obeys the following axioms:

P1. $\forall x Pxx$ (reflexivity)

P2. $\forall xy((Pxy \wedge Pyz) \rightarrow Pxz)$ (transitivity)

P3. $\forall xy((Pxy \wedge Pyx) \rightarrow x = y)$ (antisymmetry)

Events in E are never parts of each other. Rather, they are parts of objects in F . F is thus a set of composite objects that are related to members of E via P .

Each quintuple $\langle E, R, C, P, F \rangle$ is a mereological model of the non-identity theory when it obeys the following constraints:

M1. There is a fusion $f \in F$ that is composed of every event in E .

M2. There are laws on f that (i) modally constrain causal connections between events in E , dictating which events can or cannot be connected by C and that (ii) make f the right type of structure to facilitate the emergence of spatiotemporal relations.

M3. Each timelike relation in R is composed of pairs of events that are either joined by a fundamental causal relation or that are not joined by such a relation, but could be according to the laws on f .

M4. Each spacelike relation in R is composed of pairs of events that are not joined by a fundamental causal relation, and cannot be according to the laws on f .

An example of a mereological model of the non-identity theory is depicted in Figure 2.




Figure 2: *Spatiotemporal relations are composed by events. When those events are causally connected, the relation composed is timelike. When the events are not causally connected and cannot be, then the relation composed is spacelike. The left is the spatiotemporal level and the arrows represent spatiotemporal relations, the right is the fundamental level and the arrows represent causation.*

In our mereological models, spatiotemporal relations are grounded in causation in so far as those spatiotemporal relations are composed of events embedded within a causal structure. The way in which those events are embedded in a broader structure, and the laws that govern that structure together dictate the kind of spatiotemporal relations that events ultimately compose. There is still a kind of ‘grounding by absence’ but only in the sense that spacelike relations are identical to fusions of events of a particular type, a type that gets characterised negatively, in terms of the lack of causal connectedness. There is no sense in which spatiotemporal relations are somehow composed of absences.

Note that in the models we are proposing, causal relations themselves are not parts of spatiotemporal relations. The only parts that spatiotemporal relations have are events, it is just that the overall structure dictates which events compose which relations. One could expand the view somewhat, and take causal relations themselves to be parts of spatiotemporal relations. One reason for not going this way is that only some spatiotemporal relations would be able to have causal relations as parts, because only some spatiotemporal relations are grounded in causal connections that are present at the

fundamental level. Since this picture is, on the face of it, less unified, we set it aside.

Once we have a mereological model of the non-identity theory, we can then apply it to causal set theory to interpret the theory. Doing so is potentially useful as it can help us to provide a model for how the theory works. In order to fully interpret causal set theory in line with our mereological model, we need to add a minimal modification to the core structure of the theory. In addition to a fundamental causal relation, we also require a mereological relation which binds causal set elements into composite objects.

Call the resulting causal sets: causal fusions, and call an interpretation of causal set theory that includes causal fusions: a mereo-causal interpretation. Causal fusions still satisfy the two axioms of causal set theory described above (\leq is a partial order and causal fusions are locally finite). However, they also satisfy the basic mereological axioms specified for our mereological models. The dynamics still involves a growth in the elements in E related by \leq . But as new elements are added to E , those elements are related via parthood relations, and so new composites built from elements in E are added as well.

Some of the models of the mereo-causal interpretation are members of the class of models specified for the non-identity theory. In particular, any model in which the mereological relation over causal set elements produces composite objects that can then be interpreted as spatiotemporal relations will be a model of the non-identity theory. Because the mereo-causal interpretation of causal set theory preserves the core elements of causal set theory, when it comes to the physics we can expect it to be physically equivalent to standard causal set theory, and to play the same role in providing a basis for a quantum account of gravity. However, the mereo-causal interpretation takes standard causal set theory a bit further in so far as it proposes to use the causal set elements as the mereological components of spacetime relations.

The mereo-causal interpretation of causal set theory imbues that theory with greater explanatory power. For, as it stands, it is not entirely clear how spacetime emerges in causal set theory. What we

have, at the moment, are purely mathematical results that show how a causal structure can be mapped to a spatiotemporal one. Proponents of causal set theory are silent on how we should interpret this mapping, metaphysically speaking. The mereo-causal interpretation yields one option for understanding the mapping relation at issue. According to the mereo-causal interpretation, spatiotemporal relations are composed of events, where the parts of spatiotemporal relations are the causal set elements along with the fundamental relations of precedence between them. The mapping is a reflection of this deeper mereological fact.

The way that spacetime emerges in mereo-causal causal set theory, then, can be understood as the way described by the non-identity theory: it is a form of causal emergence, whereby spatiotemporal relations depend for their existence on a more fundamental causal structure. Note that, on this interpretation, it is not exactly the case that the fundamental causal relation posited within causal set theory corresponds to causation in spacetime. This might seem to be a problem, for it is sometimes stated that, in causal set theory, the fundamental causal order behaves as a spatiotemporal causal relation under certain conditions. We take this to be a point of interpretation about the theory, however, and not something delivered by the physics itself. The physics doesn't say much about whether causation at the spatiotemporal level just is the fundamental causal order \leq or whether it is merely grounded in it. Given the problems already discussed with identifying spatiotemporal structure with fundamental causal structure, the interpretation we provide seems better for exactly the reasons that the non-identity theory seems superior to the identity theory.

5.2. Causation and Interventionism

As noted at the beginning of this section, there are two interpretive questions open for the non-identity theory. The second question concerns causation. According to the non-identity theory, causation is fundamental and spacetime is not. On many metaphysical accounts of causation, however, spatiotemporal notions are used to provide a metaphysical account of what causation is. Consider, for instance,

a process-based theory of causation, of the kind advocated by Salmon (1997, 1998), Kistler (1998), and Dowe (2007). On this account, causation is analysed as the intersection of two worldlines in spacetime, across which a conserved quantity (or similar) is transferred.

Alternatively, consider Lewis's counterfactual theory of causation (Lewis, 1974). For Lewis, causation is a matter of counterfactual dependence. Whether a counterfactual of the right kind is true, however, depends, at least in part, on a similarity measure between worlds that aims to achieve a near perfect match in particular matters of fact distributed throughout spacetime (Lewis, 1979).

It would thus appear that the usual approaches to causation are unavailable options for understanding causation in the context of the non-identity theory. The question thus arises as to how causation should be understood in the context of that theory. Our view is that we should understand causation through the lens of Woodward's (2005) interventionist theory of causation. Woodward does not aim to provide a reductive account of causation in metaphysical terms. Instead, Woodward takes causation to be a primitive notion, and then seeks to elucidate causation via its conceptual connections to other important notions, namely those of manipulation and control (which underwrite experimentation in science).

That is how we understand causation in the context of the non-identity theory. Causation is a metaphysically primitive notion, and thus not one that we can analyse (and certainly cannot analyse in spatiotemporal terms). It does not follow from that, however, that there is nothing we can say to positively characterise the fundamental causal relation at issue. In order to provide such a positive characterisation of causation we can take, as our starting point, Woodward's notion of a *total cause*, which he defines as follows:

(C) X is a total cause of Y if and only if under an intervention that changes the value of X (with no other intervention occurring) there is an associated change in the value of Y.
(Woodward, 2005, p. 73)

Note that there is no mention at all of spacetime in the above definition of a 'total cause'. For this,

we need only the notion of an intervention, and the idea of a change to the value of a variable. This notion of *change* might seem to be spatiotemporal but it can be interpreted as a purely *modal* notion. The idea is that were one to intervene on the value of X , the value of Y would be different to what it is actually. Such interventions can, but need not, be changes taking place over time. The variation of the value of the variable can be viewed as a difference in the modal space or the modal dimension, rather than as a variation in the temporal dimension, or in spacetime in the temporal direction.

We couple causation in this sense to the notion of a *causal graph*. A causal graph is a directed acyclic graph consisting of nodes representing variables, and arrows which represent causal connections between nodes. Each arrow is governed by a structural equation which specifies the way in which a change in the value of one variable makes a difference to the value of a connected variable. Each node corresponds to an event, and each arrow represents causal dependence. A path through the graph is a sequence of nodes and edges. The non-identity theory thus takes the fundamental structure of reality to be modelled by an appropriate causal graph, in which events are causally linked. Because all partial orders correspond to some graph, each of our mereological models can be represented in this form.

Interventions on a variable X inside the graph can be modelled as the introduction of a new variable I —the intervention variable—that obeys the following conditions:

1. I causes X .
2. I acts as a switch for all the other variables that cause X . That is, certain values of I are such that when I attains those values, X ceases to depend upon the value of other variables that cause X and instead only depends on the value taken by I .
3. Any direct path from I to Y goes through X . That is, I does not directly cause Y and is not a cause of any causes of Y that are distinct from X except, of course, for those causes of Y , if any, that are built into the $I - X - Y$ connection itself; that is, except for (a) any causes of Y that are effects of X (i.e., variables that are causally between X and Y) and (b) any causes of Y that are

between I and X and have no effect on Y independently of X .

4. I is independent of any variable Z that causes Y and is on a directed path from I to Y that does not go through X . (Woodward, 2005, p. 75)

As is well known, Woodward's theory of causation is circular. Causation is defined in terms of interventions, and interventions are defined in terms of causation. The circularity is not problematic because the aim is not to provide a definition of causation but instead to give a functional characterisation of causation. What matters for our purposes is that none of Woodward's conditions for interventions make any reference to spatiotemporal notions. Rather, the conditions are specified entirely in terms of counterfactual relationships between variables in a causal graph, which represent some physical structure.

So long as there is a sufficiently rich counterfactual structure, it will be possible to make sense of interventions into a physical system, and thus of causation. What does it take to have counterfactual structure? All one typically needs is nomic dependence between elements. For when there is a lawful connection between some x and some y then it is usually the case that x depends on y with a degree of modal strength that can be exploited for the purposes of scientific manipulation and control. Laws typically establish a modal link between events which, in turn, mediates counterfactual dependence. This is why laws are sought after in science, and partly why we care about them in the first place. If the laws of nature need not feature spatiotemporal relations at the most fundamental level, then, there is no need for spacetime to produce structure enough for interventions and thus for causation.

This nicely pushes the question of whether there can be causation that is not fundamentally spatiotemporal back onto physics. If a physical theory can be produced that specifies lawful connections between entities that are not spatiotemporally connected, but that support interventions, then such a theory can be interpreted in line with the non-identity theory we have outlined. Ultimately, the question of whether there can be fundamental causation without spacetime just is the question of whether a viable physical theory along the relevant lines can be developed.

In the case of causal set theory, for instance, there does seem to be sufficient nomic structure to underwrite causation in Woodward's sense. We have in mind here the law of sequential growth. As discussed, this law determines the way in which causal sets 'grow' to include new elements. Ultimately, the law delivers a set of elements joined by causal relations. What the law does is specify dependencies between each of the elements in the causal set.

These dependencies underwrite interventions. For each causal set element, we can assign it a binary variable. We can represent the two states of the variable as 1 and 0 respectively. When the variable is set to 1, then the causal set element exists. When the element is set to 0, then the causal set element does not exist. Each causal set element depends on some other elements for its existence. What the law of sequential growth does is provide information about what these dependencies are. The law also tells us what would happen were we to remove an element. If we were to remove an element, then all of the elements on which it depends would be wiped out.

We can model this in a very simple way as follows. Suppose we have just three elements: a , b and c . The law of sequential growth tells us that b depends on a , and c depends on b . We can thus define a set of structural equations for the system. If we let X , Y and Z be the variables for the states of a , b and c respectively, then the structural equations for the system are just: $Y = X$, $Z = Y$, and $X = 1$. We can then model interventions on the system as changes that alter the value of specific variables. So, for instance, if we change the value of X to 0, then we wipe out b and c . Similarly, if we change the value of Y to 0 then we wipe out c . Thus, we can reason about what would happen under each of these interventions into the system.

Now, we have to be a bit careful here. Strictly speaking, Woodward's notion of an intervention requires adding a cause into the system which alters the value of a variable. This does not easily work for causal set theory. The law of sequential growth doesn't really describe what happens when we add some new element—which would be a cause on this picture—as a way of wiping out elements already in a causal set. Rather, the law just gives us a simple dependence of one element on another.

As Frisch (2014) discusses, however, it is possible to broaden Woodward's definition of an intervention so that it no-longer relies on adding an element into a causal graph as the basis for an intervention. Rather than taking an intervention to be a matter of finding some causal switch that brings about a change, we can instead just treat an intervention as nothing more than the practice of 'setting' the value of some variable inside a causal graph. This corresponds to the notion of an intervention advocated by Pearl (2000). As Frisch puts the point:

An *intervention* into the variable X_i , according to Pearl's account, then consists of removing the structural equation for X_i from the causal model and replacing it with an equation that fixes the value of X_i to some fixed x_i . That is, formally an intervention is represented by removing the equation $x_i = f_i(pa_i, u_i)$ from the model and replacing it with some x_i . Pearl calls such an intervention an "atomic intervention," which can be denoted by "*do* ($X_i = x_i$)" or "*do*(x_i)". (Frisch, 2014, p. 95)

On Pearl's picture, then, an intervention is not defined in terms of the addition of a new variable to a causal graph—the intervention variable—which is then defined to be a cause. Rather, an intervention is defined entirely via the causal graph: we simply reach into the graph and set the value of some variable directly, without adding a new cause. Using Pearl's notion of an intervention it is clear that causal set theory supports interventions in at least some sense, since we can imagine changing the state of individual causal set elements and then unfolding the consequences of this for the rest of the causal set. Causal set theory can thus be interpreted in line with our non-identity theory by using a notion of fundamental causation in a broadly interventionist sense.

One might disagree: the notion that Woodward or Pearl are interested in is just not the kind of notion that is apt to describe the world at the fundamental level. If the notion that these philosophers are interested in only applies at non-fundamental levels, then that's fine. We are prepared to admit that we are talking about a different notion, one that does apply at the fundamental level and that shares certain continuities with the notion that interests Woodward and Pearl. The issue here is similar to one

that we flagged at the beginning, namely that the kind of dependence that we use to rebuild the causal theory might not be considered as causal by everyone. By the same token, the notion that we are using may not be considered properly interventionist by everyone. We expect this is a purely verbal matter: what we call the dependence relation that underwrites our picture hardly matters. More important is whether it can provide foundation for spacetime, and we suggest that it can. What we can all agree on, perhaps, is that we are working with some modal notion of dependence. It could be mooted, then, that a better name for the view we defend here might be the modal theory of spacetime. Whatever one calls the view, however, it remains continuous with the historical project of understanding spacetime in causal terms.

6. Conclusion

We have argued that the causal theory of spacetime is a viable philosophical programme. It also finds support from contemporary physics and promises to help explain the emergence of spacetime in causal set theory. The rich diversity of distinct versions of the causal theory of spacetime has been under-appreciated by philosophers in recent times, as our focus on the non-identity theory shows. We have not tried to fully map the space of all possible causal theories of spacetime, though. Rather, we have elaborated a model based on the simple idea that causal relations could be more fundamental than spatiotemporal relations instead of being identical to them.

We take this non-identity model to be promising not only as a way to address philosophical issues like the causal indolence problem but also as a way to illuminate conceptual issues in the foundations of causal set theory (and, potentially, within any approach to quantum gravity that denies the fundamentality of spacetime). The model helps to explain the potential existence of causal systems that fail to implement a spatiotemporal profile in general and, in particular, makes conceptual room for non-spatiotemporal causal sets in an elegant way.

We doubt that our version of the non-identity theory is the only way to revive the causal approach.

Future work could be profitably directed at identifying new causal theories of spacetime, assessing their merits and comparing them with the non-identity theory discussed here. The causal theory of spacetime remains a promising view and one that deserves to be put back on the agenda within philosophy.

6. Acknowledgement

We thank the reviewers for their helpful comments, and Christian Wüthrich for helping us to better understand some aspects of causal set theory. The work of Sam Baron was supported by the Australian Research Council through a Discovery Early Career Researcher Award (DE180100414). The work of Baptiste Le Bihan was supported by the Swiss National Science Foundation.

References

- Baron, Sam and Miller, Kristie. 2015. "Causation sans time." *American Philosophical Quarterly* 52:27–40.
- Bennett, Karen. 2017. *Making Things Up*. Oxford: Oxford University Press.
- Bombelli, Luca, Lee, Jooan, Meyer, David, and Sorkin, Rafael D. 1987. "Space-time as a causal set." *Physical Review Letters* 59:521.
- Brightwell, Graham, Dowker, Fay, García, Raquel S., Henson, Joe, and Sorkin, Rafael D. 2003. "'Observables' in causal set cosmology." *Phys. Rev. D* 67:084031. doi:10.1103/PhysRevD.67.084031.
- Crowther, Karen. 2018. *Effective Spacetime: Understanding Emergence in Effective Field Theory and Quantum Gravity*. Springer.
- Dowe, Phil. 2007. *Physical Causation*. Cambridge: Cambridge University Press.

- Dowker, Fay. 2006. "Causal sets as discrete spacetime." *Contemporary Physics* 47:1–9.
- . 2013. "Introduction to Causal Sets and their Phenomenology." *General Relativity and Gravitation* 45:1651–1667.
- Earman, John. 1972. "Notes on the causal theory of time." *Synthese* 24:74–86.
- Frisch, Mathias. 2014. *Causal Reasoning in Physics*. Cambridge: Cambridge University Press.
- Grünbaum, Adolf. 1973. "The Causal Theory of Time." In Robert S. Cohen and Marx W. Wartofsky (eds.), *Philosophical Problems of Space and Time*, 179–208. Dordrecht: Springer Netherlands. ISBN 978-94-010-2622-2. doi:10.1007/978-94-010-2622-2_7.
- Hawking, Stephen W, King, A Ro, and McCarthy, PJ. 1976. "A new topology for curved space-time which incorporates the causal, differential, and conformal structures." *Journal of Mathematical Physics* 17:174–181.
- Huggett, Nick. 2014. "Skeptical notes on a physics of passage." *Annals of the New York Academy of Sciences* 1326:9–17.
- Huggett, Nick and Wüthrich, Christian. 2013. "Emergent spacetime and empirical (in)coherence." *Studies in History and Philosophy of Modern Physics* 44:276–285.
- Kistler, Max. 1998. "Reducing causality to transmission." *Erkenntnis* 48:1–25.
- Lewis, David. 1974. "Causation." *Journal of Philosophy* 70:556–567.
- . 1979. "Counterfactual Dependence and Time's Arrow." *Noûs* 13:455–476.
- Major, Seth, Rideout, David, and Surya, Sumati. 2009. "Stable homology as an indicator of manifoldlikeness in causal set theory." *Classical and Quantum Gravity* 16:175008.

- Malament, David B. 1977. "The class of continuous timelike curves determines the topology of spacetime." *Journal of Mathematical Physics* 18:1399–1404.
- Mehlberg, Henry. 1980. *Time, Causality, and the Quantum Theory: Volume One, Essay on the Causal Theory of Time*. Taylor & Francis.
- Nerlich, Graham. 1982. "Special relativity is not based on causality." *British Journal for the Philosophy of Science* 33:361–388.
- Paul, Laurie A. 2002. "Logical parts." *Noûs* 36:578–596.
- . 2012. "Building the world from its fundamental constituents." *Philosophical Studies* 158:221–256.
- Pearl, Judea. 2000. *Causality*. Cambridge: Cambridge University Press.
- Reichenbach, Hans. 1956. *The Direction of Time*. Dover Publications.
- Rideout, David and Wallden, Petros. 2009. "Emergence of spatial structure from causal sets." *Journal of Physics: Conference Series* 174:012017.
- Rideout, David P and Sorkin, Rafael D. 2000. "Classical sequential growth dynamics for causal sets." *Physical Review D* 61:024002.
- Russell, Bertrand. 1927. *The Analysis of Matter*. New York: Harcourt, Brace and Company.
- Salmon, Wesley C. 1997. "Causality and explanation: A reply to two critiques." *Philosophy of Science* 64:461–477.
- . 1998. *Causality and Explanation*. New York: Oxford University Press.
- Smart, J. J. C. 1969. "Causal theories of time." *The Monist* 53:385–395.

- van Fraassen, Bas C. 1970. *An Introduction to the Philosophy of Time and Space*. Columbia University Press.
- Woodward, James. 2005. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
- Wüthrich, Christian. 2012. “The structure of causal sets.” *Journal for General Philosophy of Science* 43:223–241.
- Wüthrich, Christian and Callender, Craig. 2016. “What becomes of a causal set?” *The British Journal for the Philosophy of Science* 68:907–925.
- Wüthrich, Christian and Huggett, Nick. 2020a. “Chapter 2: Spacetime from causality: Causal set theory.” In *Out of Nowhere*. Oxford: Oxford University Press. ArXiv preprint arXiv:2005.10873.
- . 2020b. “Chapter 3: The emergence of spacetime from causal sets.” In *Out of Nowhere*. Oxford: Oxford University Press. ArXiv preprint arXiv:2009.02951.
- Zalel, Stav. 2020. “The structure of covtree: Searching for manifestly covariant causal set dynamics.” *Classical and Quantum Gravity* 38:015001.