PAPER • OPEN ACCESS

Symmetry-breaking and zero-one laws

View the article online for updates and enhancements.

Symmetry-breaking and zero-one laws

Fay Dowker[1,2,5] and Rafael D Sorkin[1,3,4]

[1] Perimeter Institute, 31 Caroline Street North, Waterloo ON, N2L 2Y5 Canada
[2] Blackett Laboratory, Imperial College, Prince Consort Road, London SW7 2AZ, United Kingdom
[3] Raman Research Institute, C V Raman Avenue, Sadashivanagar, Bangalore– 560 080 India
[4] Department of Physics, Syracuse University, Syracuse, NY 13244-1130, United States of America

E-mail: f.dowker@imperial.ac.uk

Abstract

We offer further evidence that discreteness of the sort inherent in a causal set cannot, in and of itself, serve to break Poincaré invariance. In particular we prove that a Poisson sprinkling of Minkowski spacetime cannot endow spacetime with a distinguished spatial or temporal orientation, or with a distinguished lattice of spacetime points, or with a distinguished lattice of timelike directions (corresponding respectively to breakings of reflection-invariance, translation-invariance, and Lorentz invariance). Along the way we provide a proof from first principles of the zero-one law on which our new arguments are based.

Keywords: discreteness, symmetry breaking, zero-one law, Poisson process, causal set, quantum gravity

1. Introduction

Will a discrete structure prove to be the kinematical basis of quantum gravity and if so should we expect it to preserve the known symmetries of Minkowski spacetime, at least quasi-locally? One strand of thought has tended to answer these questions with 'yes' followed by 'no', and has held out effects like modified dispersion relations for electromagnetic waves as promising candidates for a phenomenology of spatiotemporal discreteness. In contrast we have maintained in earlier work that the type of discreteness inherent in a causal set cannot, in and of itself, serve to break Poincaré invariance. In [1] we offered informal arguments to this effect,

[5]Author to whom any correspondence should be addressed.

and then in [2] it was proved rigorously that a 'sprinkling' of Minkowski spacetime induced by a Poisson process can determine a rest frame only with zero probability.

This theorem, however, left open the possibility that a sprinkling, even if it could not remove all the symmetry of flat spacetime, could nevertheless cut it down to a proper subgroup H of the Poincaré group G. In this paper we will address that possibility, and provide further evidence against it, proving in particular that a Poisson sprinkling of Minkowski spacetime cannot induce an 'arrow of time' or a 'chirality', that it cannot break translation-symmetry by endowing spacetime with a distinguished lattice of points, and that it cannot break Lorentz symmetry by endowing spacetime with a distinguished 'lattice' of timelike directions. More generally we conjecture that a sprinkling will almost surely preserve the full group G, and we explain how one can potentially corroborate this expectation in any particular case (i.e. for any putative pattern of symmetry breaking) by combining the methods of this paper with those of [2].

Our new method herein will rely on a certain 'zero-one law' that governs invariant events in the theory of Poisson processes. To make the paper more self-contained, and also to provide a result of the requisite strength, we have chosen to prove the main zero-one theorem starting from nothing but general facts about probability measures. The resulting demonstration seems to us to be as simple as possible, and we hope that along with the proof per se, some of the definitions and lemmas that lead up to the main theorem will prove to be of independent interest.

After presenting and proving these lemmas in the next section of the paper, we prove the main theorem and then show how to apply it to exclude symmetry-breaking, first in important special cases, and then conjecturally in the general case. We also take the opportunity to reply, in an appendix, to some recent criticism of the theorems proven in [2].

For further background on these questions we refer the reader to [1, 2].

2. Preparing to prove a zero-one law

In the next section, we will prove a 'zero-one law' about Poisson processes, from which will follow the desired theorems on symmetry-preservation in many, if not all, cases of interest. In fact, a version of this result can be found in [3], but that theorem would not let us rule out certain important cases of symmetry-breaking. For example it would not let us exclude that a sprinkling might break the group of all translations down to a discrete subgroup, as happens for example when a liquid crystallizes. For this reason, we have decided to demonstrate *ab initio* the zero-one law we will be appealing to herein. We hope also that our development will help to clarify how and why such laws arise. In preparation, let's first review some definitions and known results from [2, 4].

Let μ be the measure that, mathematically speaking, defines our sprinkling process, which we take to be a Poisson process in $\mathbb{M}^n$, the Minkowski space of dimension n. An individual sprinkling in $\mathbb{M}^n$ is almost surely a locally finite subset of $\mathbb{M}^n$. The space of all such subsets, which we will denote by Ω, is the *sample space* of the Poisson process. A measurable subset of Ω will be called an *event*, as is customary for stochastic processes. The set of all events forms a σ-algebra that we will call the *event-algebra* $\mathfrak{A}$.

The concept of a *bounded event* will be important for our proof. By definition such an event will be one that pertains to a bounded (say compact) subset of $\mathbb{M}^n$, by which we mean more precisely the following. Let $\omega \in \Omega$ be any sprinkling, and B a subset of $\mathbb{M}^n$. We say that an event A is 'an event within B' (or is 'supported within B') if in order to know whether $\omega \in A$ it suffices to know the subset, $\omega \cap B$, of sprinkled points that fall within B. For example the event, 'There are more than five sprinkled points in B', is an event within B. We call an event bounded if it is an event within B for some bounded spacetime region B.

We will write $\mathfrak{A}_0$ for the set of all bounded events. It is not a σ-algebra, but it is still a Boolean algebra, meaning it is closed under the operations of Boolean sum and Boolean product, as

defined below. Equivalently it is closed under union, intersection, and set-difference.

It will be important for our proof that every event $A \in \mathfrak{A}$ can be built up as a (countable) logical combination of bounded events. Formally, this says that the full event-algebra $\mathfrak{A}$ is generated qua σ-algebra by $\mathfrak{A}_0$. (This basic fact about Poisson processes results directly from the way in which they are defined [3, 5].) We claim (and will shortly prove) that as a consequence, every event in $\mathfrak{A}$ is in a well-defined sense a limit of bounded events.

Before turning to the proof, we need to establish a few more definitions and some notation and lemmas. Most of the lemmas are either well known or easy to prove, but we include them for completeness, and because some of our definitions are not quite the usual ones.

Notation.   Let A and B be events. Their *Boolean sum*, $A + B$, is their 'symmetric difference', $(A \cup B) \backslash (A \cap B)$. Their *Boolean product*, AB, is their intersection, $A \cap B$.

This little-used but convenient notation exhibits explicitly that the events form an algebra over $\mathbb{Z}_2$, with identity 1 equal to the event Ω. The complement of an event A can thus be written as $1 + A$.

Definition ('Distance' between two events). $d(A, B) = \mu(A + B)$

Definition.   Let $A, A_1, A_2, A_3 \ldots$ be events in $\mathfrak{A}$. Then $A_k \to A$ means that $d(A_k, A) \to 0$.

We will also say in this situation that A is a limit of the A_k. (Our definition of distance satisfies the triangle inequality and therefore defines a metric on $\mathfrak{A}$, or more precisely a pseudometric, because two events that differ by a set of measure zero will be at zero distance without actually being equal. This in turn is why we have written that A is 'a limit' and not 'the limit' of the A_k.)

The next two lemmas will verify the triangle-inequality for d. The latter is not technically a metric, however, because $d(A, B) = 0$ does not imply that $A = B$.

Lemma 1. $\mu(A + B) \leqslant \mu(A) + \mu(B)$

Proof. $A + B \subseteq A \cup B \Rightarrow \mu(A + B) \leqslant \mu(A \cup B) \leqslant \mu(A) + \mu(B)$. $\square$

Lemma 2 (Triangle inequality). $d(A, C) \leqslant d(A, B) + d(B, C)$

Proof. $A + C = (A + B) + (B + C)$ because $B + B = 2B = 0$. Hence, in light of the previous lemma, $\mu(A + C) \leqslant \mu(A + B) + \mu(B + C)$. $\square$

Lemma 3. $|\mu(A) - \mu(B)| \leqslant \mu(A + B)$

Proof.   A Venn diagram makes this clear. More computationally, we have, since the measure μ is additive, $\mu(A) = \mu(A \backslash B) + \mu(AB)$, and similarly $\mu(B) = \mu(B \backslash A) + \mu(AB)$, whence $\mu(A) - \mu(B) = \mu(A \backslash B) - \mu(B \backslash A) \leqslant \mu(A \backslash B) + \mu(B \backslash A) = \mu(A + B)$, and similarly $\mu(B) - \mu(A) \leqslant \mu(A + B)$. $\square$

From this last lemma follows immediately the continuity of μ with respect to d, as well as that of addition and multiplication.

Lemma 4. $A_j \to A \Rightarrow \mu(A_j) \to \mu(A)$

Lemma 5. $A_j \to A$ *and* $B_j \to B \Rightarrow A_j B_j \to AB$ *and* $A_j + B_j \to A + B$
   (*In other words limit preserves Boolean sum and product.*)

Proof.   First notice that $A_j B_j + AB = A_j(B_j + B) + (A_j + A)B$, and that $A_j(B_j + B) \subseteq (B_j + B)$, while $(A_j + A)B \subseteq A_j + A$. Therefore $d(A_j B_j, AB) = \mu(A_j B_j + AB) \leqslant \mu(B_j + B) + \mu(A_j + A) = d(B_j, B) + d(A_j, A) \to 0$. The proof for $A + B$ is similar but simpler. Start with the

trivial equation, $(A_j + B_j) + (A + B) = (A_j + A) + (B_j + B)$ and apply μ to both sides. The result is $d(A_j + B_j, A + B) = \mu[(A_j + A) + (B_j + B)] \leqslant \mu(A_j + A) + \mu(B_j + B) = d(A_j, A) + d(B_j, B) \to 0$. □

Remark. We could prove in the same way that limit preserves complementation: $A_j \to A \Rightarrow 1 + A_j \to 1 + A$, but it follows already from the lemma.

The next lemma holds for any Boolean algebra of events and the σ-algebra it generates.

Lemma 6. *Every event in $\mathfrak{A}$ is the limit of a sequence of events in $\mathfrak{A}_0$*

Proof. Let $\bar{\mathfrak{A}}_0$ be the set of all such limits. Because a σ-algebra can be defined as a Boolean algebra of sets which is complete in the sense that it is closed under forming the union of an increasing sequence of sets[6], and because the σ-algebra generated by any family $\mathfrak{F}$ of events is by definition the smallest σ-algebra that includes $\mathfrak{F}$, it suffices to prove that $\bar{\mathfrak{A}}_0$ is closed under Boolean addition and multiplication, and that forming the union of an increasing sequence members of $\bar{\mathfrak{A}}_0$ does not lead out of $\bar{\mathfrak{A}}_0$ either. Since closure under the Boolean operations is the content of the preceding lemma, we only need to demonstrate closure under nested countable union. To that end, let $A = \bigcup_j A^j$ be the union of an increasing sequence of events $A^j \in \bar{\mathfrak{A}}_0$, each of which is the limit of a sequence of events A_k^j in $\mathfrak{A}_0$. It is a basic[7] result of measure theory (sometimes called 'continuity') that in this situation, $\mu(A \backslash A^j) \to 0$. But because $A^j \subseteq A$, $A + A^j = A \backslash A^j$, and we have $d(A^j, A) = \mu(A + A^j) = \mu(A \backslash A^j) \to 0$. Now choose $\varepsilon > 0$ and find an A^j such that $d(A^j, A) < \varepsilon/2$, finding next an index k such that $d(A_k^j, A^j) < \varepsilon/2$. Together, these imply that $d(A_k^j, A) \leqslant d(A_k^j, A^j) + d(A^j, A) \leqslant \varepsilon/2 + \varepsilon/2 = \varepsilon$, whence $\mathfrak{A}_0$ contains events arbitrarily close to A, as required. □

The proof of our zero-one law will rest on the previous lemma together with the following one.

Lemma 7. *If events A and B are limits of sequences of events A_j and B_j respectively, and if for each index j, A_j is stochastically independent of B_j, then A and B are also stochastically independent.*

Proof. By definition, stochastic independence of A and B signifies that $\mu(AB) = \mu(A)\mu(B)$, which accordingly is what we want to prove. But by hypothesis, we have $\mu(A_j B_j) = \mu(A_j)\mu(B_j)$. Appealing now to an earlier lemma, we can conclude from $A_j \to A$ that $\mu(A_j) \to \mu(A)$ and similarly $\mu(B_j) \to \mu(B)$, whence $\mu(A_j)\mu(B_j) \to \mu(A)\mu(B)$. On the other hand, $A_j B_j \to AB$, whence $\mu(A_j B_j) \to \mu(AB)$, completing the proof. □

3. A zero-one law and its proof

Let us say that an event $A \in \mathfrak{A}$ is *deterministic* if its probability $\mu(A)$ is either 0 or 1, but nothing in between. One also says that A obeys a 'zero-one law'. If A is a deterministic event, then either it or its complement, $1 + A$, is forbidden. In the jargon of probability theory, an event forbidden in this way 'almost surely will not happen', while its complement 'almost surely will'.

Consider now some event A, let G be the Poincaré group, and let $g \in G$ act on A by acting on the individual sprinklings ω that comprise it: $gA = \{g\omega | \omega \in A\}$. By the *invariance group* of A we mean the subset H of G whose elements leave A unchanged.

---

[6] Increasing means that $A^1 \subseteq A^2 \subseteq A^3 \cdots$.
[7] Basic but quite simple to prove from the axioms for a measure [4].

Theorem.　*If the invariance group of an event A contains at least one non-zero spacetime translation then A is a deterministic event with respect to the Poisson process in* $\mathbb{M}^n$.

Proof.　Observe to begin with that if the invariance group H contains the translation T, it automatically contains all powers of T; it therefore contains arbitrarily large translations. It follows for any bounded spacetime region K that H contains a translation T for which K and TK are disjoint. Now let B be an event within the bounded region K, and choose $T \in H$ so that K and $K' = TK$ are disjoint, and let $B' = TB$. Since B is an event within K and B' is an event within K', and since K is disjoint from K', B will be stochastically independent of B', this being a basic feature of Poisson processes.

　　Now let A_k be a sequence of bounded events such that $A_k \rightarrow A$. Such a sequence exists by lemma 6. We have just seen that for each index k, there is a translation $T_k \in H$ such that A_k and $A'_k = T_k A_k$ are stochastically independent.

　　Moreover, we claim (and this is the key to the proof) that these translated events A'_k *also* converge to A. To see why, recall first that by the definition of H, the event A is not altered by any of the T_k, i.e. $T_k A = A$. Then since the Poisson-process measure μ is itself translationally invariant, we have $d(A'_k, A) = d(T_k A_k, A) = d(T_k A_k, T_k A) = d(A_k, A) \rightarrow 0$ as claimed.

　　We now have two convergent sequences of events whose individual terms are stochastically independent. According to lemma 7, this entails that the limit-events are also stochastically independent. But we just proved that these limit-events are both equal to A, whence A is independent of itself! As an equation, this says that $\mu(AA) = \mu(A)\mu(A)$, or $\mu(A) = \mu(A)^2$, since of course $AA = A$. The only solutions of this equation being $\mu(A) = 0$ or $\mu(A) = 1$, the theorem is established.　□

4. Can a sprinkling break Poincaré invariance?

The theorem just proven will let us demonstrate several results that rule out in various cases that a sprinkling can break one of the symmetries of $\mathbb{M}^n$. When combined with the analogous results from [2], we expect that all cases of physical interest will be spoken for. To make this plausible we now apply our theorem to some prototypical examples.

4.1. A sprinkling cannot determine an orientation

As a first example let's ask whether a Poisson sprinkling can break one of the reflection-invariances by favoring either a particular spatial or temporal orientation, or a particular overall orientation. The reasoning being the same in all these cases, let's take for definiteness the case of an overall orientation (which is preserved by CPT but not CP or T). The question is then, can a sprinkling—an individual realization of the Poisson process—determine (with non-zero probability) a specific orientation O?

　　Of course only two orientations are possible, say O_1 and O_2, so our question reduces to asking for the probability p that the sprinkling will favor O_1 over O_2. By symmetry of the Poisson process, p is also the probability that it will favor O_2 over O_1. For maximum generality, we also admit that it might favor neither, so that p might be strictly less than $1/2$. We claim in fact that $p = 0$.

　　To prove this consider the event A that the realization (call it ω) favors O_1. Since an orientation can be thought of as an equivalence class of orthonormal tetrads (if $n = 4$), and since an orientation is something global, the tetrads are located nowhere in particular (or if you like they are located everywhere). The event A is thus trivially invariant under all translations. (If ω determines O and if T is any spacetime symmetry, then $T\omega$ must determine TO, which as we just saw, is O itself when T is a translation.)

Our theorem then informs us that A is a deterministic event, whence either $p = 0$ or $p = 1$. But since $p \leqslant 1/2$ in any case, the only consistent possibility is that $p = 0$, as claimed. Thus, a sprinkling will almost surely leave the reflections unbroken.

One might wonder whether something would go wrong here if the sprinkling determined *more* than just an orientation. What if it also determined a distinguished location in space-time, for example? In fact nothing would go wrong because we assumed nothing about what else ω might be able to determine. The event A would still be defined and would still be translation-invariant because it would gather together all the ω which favor O_1 irrespective of which location they might also favor.

On the other hand, the doubt we have just sought to dispel does point to a perennially confusing ambiguity that lurks in a phrase like 'A sprinkling cannot break T-reversal'. Is it saying that the particular isomorphism $t \rightarrow -t$ is (in some coordinate system) a symmetry (meaning in the present context that it belongs to the invariance group H) or is it only saying that a sprinkling cannot prefer a direction of time? The difference shows up famously in discussions of the standard model of high energy physics, where people are wont to say that time-reversal is broken but that the laws of physics introduce no arrow of time because CPT is a symmetry that reverses any putative arrow. *What our proofs in this paper establish directly is the second kind of statement, which only indirectly bears on the first.*

4.2. A sprinkling cannot break translation-symmetry by determining a spacetime lattice

In the orientation example we just treated, the tetrads acted as a kind of order-parameter or Higgs field responsible for the (putative) symmetry breaking. We take it as an article of faith that this will always be the case: if a sprinkling breaks a spacetime symmetry it will be because one can deduce from it some geometrical object X whose invariance group H is a proper subset of the full group G of symmetries. (In the case of Minkowski spacetime, which is our main interest, G will be the Poincaré group including all of its connected components. In the case of Euclidean space, to which our analysis also applies, G will be the Euclidean group, etc.)

In the present subsection, we ask whether a Poisson sprinkling can break the translation symmetry of spacetime. For this to happen, X would have to be for example a distinguished 'origin' in spacetime, resulting in a trivial H of no residual symmetry. This possibility was ruled out in [2], but X could also be a rectangular lattice of spacetime points, resulting in an H identifiable with the subgroup of translations that preserve the lattice. (This situation is familiar from crystallization, and 'crystal group' might be an apt name for H. As this name suggests, the full H might include some rotations, etc, but we will ignore them here since our concern in this example is just with translations. Thus we will for now limit G just to the translations.)

Suppose now that some sprinkling ω determines the lattice L. Reasoning as before from the overall G-invariance of the Poisson *process*, we see that other sprinklings must be able to determine other lattices, all of them equally probable. The lattices obtainable in this manner can, in the familiar way, be identified with the elements of the coset space G/H (topologically a torus).

Fix now a particular lattice L_1, and let p be the probability that L_1 will result from a sprinkling. Or more correctly (since we do not want p to vanish trivially), introduce a small rectangular neighborhood $\tilde{L}_1$ of L_1 and let A be the event: 'The sprinkling ω determines a lattice L belonging to $\tilde{L}_1$'. If the neighborhood $\tilde{L}_1$ was chosen suitably, A will be invariant under H, the invariance group of L_1, and we define $p = \mu(A)$.

The event A is the analog of the event of the same name in the orientation example, and from here onward, we can proceed exactly as before. On one hand, since H contains nontriv-

ial translations, A is deterministic, thanks to our theorem[8]. On the other hand, $p = \mu(A) < 1$ because there are other 'fuzzy-lattice events' which are just as probable as A is with respect to our Poisson process. Therefore $p = 0$ is the only possibility, and a sprinkling will almost surely leave the translations unbroken.

Remark. Exactly the same argument goes through for lattices L in Euclidean space.

4.3. A sprinkling cannot prefer a timelike direction: two methods of proof

This was the main theorem proven in [2] by a different method that assumed only that the sprinkling process was invariant under Lorentz transformations. In this paper, we are assuming more specifically that our sprinkling process is a Poisson process. To what extent this is a loss of generality is unclear, since at present there seems to be no known example of a sprinkling process that is Poincaré invariant without actually being Poisson (barring the trivial exception of a convex combination of Poisson processes of different densities).

Let us compare and contrast the two methods of proof.

Following the pattern established with the previous two examples, suppose that a sprinkling ω could determine the timelike unit vector u. Let G be the Poincaré group, as before, and let $H \subseteq G$ be the subgroup that acts as the identity on u. (Thus H contains the translations together with the rotations and spatial reflection associated with the timelike direction u.) The quotient G/H can then be identified with the (two-sheeted) unit hyperboloid in $\mathbb{M}^n$. Consider as before the sprinkling-induced correspondence $\omega \to u$ and express it as a partial function $F : \Omega \to G/H$ (it is partial because its domain might not be all of Ω). Continuing to reason as before, we learn that F induces on G/H a (subnormalized) probability distribution ν. Because it must be invariant under G, we know also that ν could only be a constant density on G/H.

At this point the two methods part ways. The method of [2] simply notices that unless $\nu = 0$ its integral over all of G/H would be infinite, whereas in fact it cannot exceed unity (being subnormalized). The only way out of this contradiction is that the domain of F is a measure-zero subset of Ω. The method of this paper, on the other hand, reaches the same conclusion by introducing a bounded subset S of G/H and observing that the event A given by '$F(\omega) \in S$' is translation invariant since u is a global object, like the orientations in our first example. Hence A is deterministic, and $\nu(S) = \mu(A)$ can only be 0 or 1, whence it must be 0 since it cannot be 1.

How then do the two methods differ? Both proceed from the same uniform density ν on G/H, but they presuppose different things about H and G/H. The first method lives off the fact that G/H has an infinite volume. The second lives off the fact that H contains a nontrivial translation. Thus, the first method works when H is 'sufficiently small', the second works when H is 'sufficiently big' (but not so big that G/H fails to contain at least two points. In that case $H = G$ and there is no breaking at all.)

In the previous two examples, the first method would not have worked because G/H was compact and hence of finite volume. On the other hand the second method would have trouble if the sprinkling were trying to break translation-invariance completely by picking out a unique favored point or 'origin'; in that case H would contain no translations. We would conjecture that in all cases of interest at least one of the two methods will work. This would be true, for example, if G/H necessarily had infinite volume whenever H failed to contain a translation.

---

[8] In the previous example the full strength of our theorem was not needed, because H there included the entire translation group.

4.4. A sprinkling cannot prefer a 'lattice' of timelike directions

As a last illustration of the second method, let us consider the possibility that 'X' is not a single timelike direction but an infinite set of them which is invariant under a discrete subgroup of the Lorentz group G[9]. It might seem surprising that such a subgroup exists at all, but many instances are known. One of the most interesting is comprised of the set of Lorentz transformations that leave invariant the integer lattice $\mathbb{Z}^4$ in $\mathbb{M}^4$ [6, 7]. The elements of X itself can then be taken to be the unit vectors pointing from the origin to the points of L. Let us focus on this example.

It seems that there are general theorems of algebraic geometry which imply in this case that orbit of such an X under the action of the Lorentz group, though not actually compact, has only a finite volume [8]. Our first method of proof would then not apply. The second method does apply however for the same reason it applied to a single timelike direction, our X's being by definition translation invariant.

5. What does it all mean?

We don't have access to all of spacetime, and in any case we don't live in $\mathbb{M}^4$. What then is the physical relevance of theorems about sprinklings of a flat spacetime? Recall that the sprinkling of a Lorentzian manifold M has only a kinematical and not a dynamical significance. It is meant to provide a causal set typifying those that could be the substructure of M[10]. If in this paper we have taken M to be literally $\mathbb{M}^4$, this is only an idealization of some approximately flat region R within the larger Universe. What we'd really like, then, is not only a global proof of Poincaré invariance, but a quasilocal result that would quantify how much anisotropy or inhomogeneity remains, depending on the size of R. Our rigorously proven theorems are but a first step toward such an analysis. (As usual there's a trade-off between beautiful theorems and applicability!)

In Euclidean space, such an analysis seems near at hand. To each spatial point x we can associate the line that passes through it and its nearest sprinkled neighbor. Each such line breaks the rotation symmetry *at that point* to $\mathbb{Z}_2$, which is of course why rotations cannot literally be a symmetry of a sprinkling but only so in an average sense. It is equally clear, though, that these lines fluctuate wildly in direction, so the anisotropy dies out rapidly with the size of the region one considers. Similarly, one would expect any localized inhomogeneities to wash out on larger scales so that translation-invariance would return.

In Minkowski space something similar is plausibly true, but in relation to the Lorentz subgroup of the Poincaré group, there's a complication; both the size *and* the shape of the region R are important. Nevertheless we would still expect to get a rapidly fluctuating array of lines that are, in the natural rest-frame of the region[11], nearly null, and so the breaking would again die out rapidly as R grew. Only now in a finite region we won't restore all of the Lorentz group, but only those boosts that are small enough for R to accommodate. This 'boundary effect' (or 'shape effect') has no analog in the Euclidean case, but otherwise the two situations seem quite similar.

---

[9] To be mathematically impeccable, we should point out that G here is not literally a subgroup of the Poincaré group, but of its quotient by the translations. That is, G doesn't act on spacetime itself, which is strictly speaking an affine space, but rather on the associated vector-space.

[10] Even this statement ignores that quantum spacetime is expected to be more like a 'superposition' of causal sets than a single one. Moreover, we only expect a sprinkling to be a good model after a certain amount of coarse-graining, e.g. if at small scales the structure of spacetime were of Kaluza–Klein type.

[11] What is the 'frame of the region'? Well, find two points x, y in R such that the order-interval $I(x, y)$ has the biggest volume possible. The line through x and y then defines the rest-frame in question. Some such prescription ought to be adequate in most cases.

Beyond these kinematic questions of global theorems vs quasilocal applicability, what we ultimately care about are consequences for the dynamics. Would a massless scalar field living on a Poisson sprinkling propagate via a modified dispersion relation, as has been suggested for discrete structures? The answer depends obviously on how the dynamics is formulated, so it is impossible to answer categorically. But our theorems are significant precisely because they indicate that the answer will be 'no'. (We ignore here the possibility of dynamical spontaneous symmetry breakings which have nothing to do with kinematical discreteness.)

Which doesn't mean there might not be *other* 'dispersive' or diffusive effects consistent with all the spacetime symmetries. We hope that there are, because they would be highly constrained by the symmetry, and would potentially provide phenomenological evidence of discreteness! [1, 9] Indeed, such effects, although not yet seen experimentally or observationally, have already begun to be studied in extant theories that describe the dynamics of particles and/or fields on a background causal set. (For examples of such theories, see [10])

But even these reflexions are not the end of the story. Beyond dynamics on a fixed, background causal set, we need ultimately to understand the effects of the causal set itself being dynamical (i.e. of quantum gravity). Our theorems here are merely a first indication of how things are likely to turn out.

Acknowledgments

Appendix A. Reply to Kent [11]

In a recent paper [11], Kent has disputed our interpretation of the theorems proven in [2]. As far as we can see, he puts forward three main criticisms, and we take this opportunity to explain why we think they are unfounded. We hope also, that our comments will help bring into focus the conceptual background to both the work in [2] and its extension here.

Kent's primary complaint seems to be that attention should fall on what he calls 'sprinklable sets' instead of sprinklings, where a sprinklable set is an isometry equivalence class of sprinklings. This amounts to treating Poincaré symmetries as if they were merely gauge, contrary to the way most physicists understand them. (We follow here the widely used terminology that draws a distinction between 'gauge transformations' that, like coordinate transformations, merely alter the description without affecting physical reality, and 'symmetries' which effect genuine physical changes. It is, for example, because one treats translations as symmetries that it is meaningful to speak of the energy-momentum vector of a system.) We believe that the majority viewpoint is in this case the appropriate one. As highlighted earlier, we don't live in $\mathbb{M}^4$ but in a cosmos that is highly curved on large scales and near to black holes, etc. In such a Universe a flat spacetime can only be an idealization of a nearly flat local region R. But as soon as you remember that all such regions exist within an enveloping spacetime, you realize [12] that local translations, rotations, and Lorentz-boosts are in the larger context *not* pure gauge, because they move a subsystem around relative to its environment. They are rather real

physical changes idealized as what one might term 'partial gauge transformations'[12]; and one really ought to think of $\mathbb{M}^4$ as being referred to an 'external frame'—a laboratory, the fixed stars, etc. (If the whole of spacetime really were $\mathbb{M}^4$, one might have to rethink the status of the Poincaré group, but obviously that is not the case.) Thus sprinklings and not sprinklable sets are the appropriate objects of study.

Having replaced sprinklings by sprinklable sets, Kent then argues, if we understand him, that the zero-one law that holds for propositions about sprinklable sets is a bad thing because it means in some sense that one cannot say anything interesting about a sprinklable set created by a Poisson process. Of course, this criticism cannot be sustained if, as we have just argued, it is sprinklings and not sprinklable sets that are physically relevant. But instead of just stopping with this comment, perhaps we should add that (as explained by Kent himself under the heading 'A lacuna in the BHS theorem') a question like 'Does the sprinkling determine a timelike direction?', still makes sense as a question about sprinklable sets. The corresponding event in the sample-space Ω_S of sprinklable sets is simply the union of all the events in $\mathfrak{A}$ that belong to specific timelike directions; and it still has measure zero. (See [14] for how Ω_S is related to Ω.) Since this question and others like it hold the keys to deciding whether a sprinkling can break a spacetime symmetry, we cannot agree that the σ-algebra of Ω_S is too sparse to contain events of physical interest, even if one chooses to study it instead of Ω.

But independent of 'sprinkling vs sprinklable', could it be that something else is behind the criticism? There are hints in [11] that one is thinking of the Poisson process as a kind of dynamical theory of causal sets. If one were to think of it in this way, then one might feel uncomfortable that every event in this theory would be deterministic. For some purposes that might be an interesting observation, but it is in any case not relevant to causal set dynamics. As described in the previous section, sprinklings within causal set theory play only the kinematical role of helping to define the relationship between a causal set and the corresponding spacetime continuum. Dynamical laws ('laws of motion') meant for causal sets can presuppose no background spacetime, and are envisioned as defining a stochastic process of growth which, as it were, builds up an evolving causal set element by element.

Remark. Suppose that in some context one actually did want to interpret the Poisson process as a discrete dynamics for Minkowski spacetime. There is only one $\mathbb{M}^4$-geometry, and since every question you can ask about its structure thereby has a unique yes-or-no answer, would not a zero-one law for such questions be exactly what you would want? It would suggest that your dynamics had reproduced $\mathbb{M}^4$ as well as it could, consistent with discreteness.

Kent's third criticism seems to be that reference [2] proved the wrong thing, or at least failed to prove some things it needed to prove. In effect he has brought forward a new requirement that anyone claiming to establish Poincaré invariance needs to satisfy, which he states as follows. 'One needs to show that, given any data that leave some continuous subgroup of the Lorentz group as a symmetry in the continuous case, there is no mathematical construction that breaks this symmetry in the discrete case'.

To see what this means, consider for simplicity the Euclidean question whether a sprinkling can prefer a spatial direction, thereby breaking isotropy. This was a question that could not be answered in [2], but which we have answered in the negative in the present paper.

---

[12] By partial gauge transformation we mean an operation which is locally indistinguishable from a gauge transformation but which only acts nontrivially on a subsystem or region while leaving the surroundings unchanged. Most if not all symmetries can be understood as partial gauge transformations. See for example the brief discussion of this concept (though not under this name) in section 1 of [13].

Now consider the different question whether a sprinkling could determine a spatial direction if one provided in addition a marked spatial point or 'origin'. As pointed out in [2], the answer to this question is 'yes'. Does this constitute a breaking of isotropy? Kent thinks it does, whereas we think it does not, because the required extra information is in reality absent[13]. We therefore disagree that there is some kind of 'lacuna' in the theorems of [2] or this paper. For us the most pertinent questions are the *intrinsic* ones, that ask whether a sprinkling *in and of itself* can break a symmetry.

The above is of course not meant to claim that the theorems in [2] settled every question one might want to ask. On the contrary, our concern in this paper has been to complement those theorems by analyzing a larger class of symmetry-breaking scenarios than was possible with the tools of [2] alone. And beyond that loom the whole series of questions adumbrated in the previous section.

It is in connection with the latter questions that Kent's 'extra information' might become relevant. He invokes for example a particle moving through a medium of sprinkled points (in $\mathbb{M}^4$, but let's stay Euclidean for convenience). The particle itself 'marks a point', and so it can in fact see some anisotropy. It will then swerve from a straight line, and this effect could be noticed. Very good! This is precisely the type of effect one expects from discreteness. But what's important is the inference that—precisely because isotropy is intrinsically *preserved*—the diffusion equation describing these swerves will be rotationally invariant. Just such an equation, in its Lorentzian guise, was brought forth in [1] as a possible phenomenological manifestation of an underlying causal set. The extrinsic information provided microscopically does something observable, but in a manner that respects the intrinsic global symmetry.

Remark.　Apropos of Kent's remarks on local Lorentz invariance, we have noticed that certain passages in [2] could lead readers to interpret that ambiguous phrase in a manner less like what it would mean in the context of this paper, and more like what it means in connection with fields of orthonormal tetrads. If so, we hope that the reflections in the previous section concerning what one might call 'local Poincaré-invariance' (which, be it noted, includes translations) will have made it clear that the words local or quasilocal are in the present context not meant to point to any extrinsically given location or marked point in spacetime; they are meant rather to evoke the kind of approximately flat region R expounded on above under the heading 'What does it all mean?'

ORCID iDs

Fay Dowker ● https://orcid.org/0000-0002-6652-1058

References

[1] Dowker F, Henson J and Sorkin R D 2004 Quantum gravity phenomenology, Lorentz invariance and discreteness *Mod. Phys. Lett.* A **19** 1829–40
[2] Bombelli L, Henson J and Sorkin R D 2009 Discreteness without symmetry breaking: a theorem *Mod. Phys. Lett.* A **24** 2579–87
[3] Last G and Penrose M 2017 *Lectures on the Poisson Process* vol 7 (Cambridge: Cambridge University Press)

[13] In the Lorentzian example considered in [11], the extra information is a timelike direction, but one is still asking about spatial rotations.

[4] Kolmogorov A N and Fomin S V 1961 *Measure, Lebesgue Integrals, and Hilbert Space* (New York: Academic) (translated by Natascha Artin Brunswick and Alan Jeffrey)

[5] Kingman J F C 1993 *Poisson Processes, Oxford Studies in Probability* vol 3 (Oxford: Clarendon)

[6] Schild A 1949 Discrete space-time and integral Lorentz transformations *Can. J. Math.* **1** 29–47

[7] Saravani M and Aslanbeigi S 2014 On the causal set-continuum correspondence *Class. Quantum Grav.* **31** 205013

[8] Avner Ash Personal Communication

[9] Philpott L, Dowker F and Sorkin R D 2009 Energy-momentum diffusion from spacetime discreteness *Phys. Rev.* D **79** 124047

[10] Sorkin R D 2011 Scalar field theory on a causal set in histories form *J. Phys.: Conf. Ser.* **306** 012017

Dable-Heath E, Fewster C J, Rejzner K and Woods N 2020 Algebraic classical and quantum field theory on causal sets *Phys. Rev.* D **101** 065013

Saravani M, Sorkin R D and Yazdi Y K 2014 Spacetime entanglement entropy in $1 + 1$ dimensions *Class. Quantum Grav.* **31** 214006

Sorkin R D and Yazdi Y K 2018 Entanglement entropy in causal set theory *Class. Quantum Grav.* **35** 074004

Zwane N Cosmological tests of causal set phenomenology *PhD Thesis* University of Waterloo

Dowker F, Henson J and Sorkin R D 2010 Discreteness and the direct transmission of light from distant sources *Phys. Rev.* D **82** 104048

Sorkin R D 2009 Does locality fail at intermediate length-scales? *Approaches to Quantum Gravity—Towards a New Understanding of Space and Time* ed D Oriti (Cambridge: Cambridge University Press) pp 26–43

[11] Kent A 2018 Are there testable discrete Poincaré invariant physical theories? *Class. Quantum Grav.* **35** 195001

[12] Poincaré H 1963 *Mathematics and Science: Last Essays* (New York, NY: Dover) pp 20–1 (translated from the French by John W Bolduc)

[13] Friedman J L and Sorkin R D 1983 Statistics of Yang–Mills solitons *Commun. Math. Phys.* **89** 501–21

[14] Brightwell G, Dowker F, García R S, Henson J and Sorkin R D 2003 'Observables' in causal set cosmology *Phys. Rev.* D **67** 084031