Math. Struct. in Comp. Science (2012), *vol.* 22, *pp.* 816–852. Published by Cambridge University Press 2012
doi:10.1017/S0960129511000545

Toward a fundamental theorem of quantal measure theory

RAFAEL D. SORKIN[†]

Perimeter Institute, 31 Caroline Street North, Waterloo ON, N2L 2Y5, Canada
and
Department of Physics, Syracuse University, Syracuse, NY 13244-1130, U.S.A.
Email: rsorkin@perimeterinstitute.ca

Received 3 April 2011

In this paper we address the extension problem for quantal measures of path-integral type, concentrating on two cases: sequential growth of causal sets and a particle moving on the finite lattice $\mathbb{Z}_n$. In both cases, the dynamics can be coded into a vector-valued measure μ on Ω, the space of all histories. Initially, μ is just defined on special subsets of Ω called cylinder events, and we would like to extend it to a larger family of subsets (events) in analogy to the way this is done in the classical theory of stochastic processes. Since quantally μ is generally not of bounded variation, a new method is required. We propose a method that defines the measure of an event by means of a sequence of simpler events that in a suitable sense converges to the event whose measure we are seeking to define. To this end, we introduce canonical sequences approximating certain events, and we propose a measure-based criterion for the convergence of such sequences. Applying the method, we encounter a simple event whose measure is zero classically but non-zero quantally.

1. Introduction

In order to define area, even for something as simple as a disk of unit radius, we need to invoke an extension theorem. In a systematic development (Kolmogorov and Fomin 1961) of plane measure, we begin by defining the measure μ of an arbitrary rectangle, and then seek to extend the set-function μ unambiguously to subsets of the plane that can be made from rectangles through countable processes of union and complementation (these sets comprising the σ-algebra generated by the rectangles). The unit disk is such a subset, and (if we take it to be open) it can obviously be built up as the disjoint union of a countable family of rectangles. But this can be done in an infinite number of different ways, and we need to know that the net area of the rectangles is always the same, no matter which decomposition we choose and no matter what order we choose to perform the resulting sum. The theorem that guarantees this consistency is known as the Kolmogorov–Carathéodory extension theorem, but it might also be called

[†] Research at the Perimeter Institute for Theoretical Physics is supported in part by the Government of Canada through NSERC and by the Province of Ontario through MRI.

the 'fundamental theorem of classical measure theory'. Not only is it used to construct Lebesgue measure, but it plays a central role in defining stochastic processes like the Wiener process, which is a mathematical model of Brownian motion that also describes the Wick rotated path-integral for a non-relativistic free particle on the line.

In this sort of application, we are dealing with a probability measure on a space of paths or, more generally, 'histories', and the possible values of μ are therefore positive real numbers between 0 and 1. However, when we try to define a genuine path integral in real time (as opposed to Wick-rotated, imaginary time), we encounter complex amplitudes that can be arbitrarily large and of any phase. Once again, there are some especially simple sets of paths, called 'cylinder sets', which are analogues of the rectangles, and from which the more general sets of interest can be built up. However, the sums that arise in this case no longer converge absolutely. In technical terms, the complex measure we are trying to extend is not of bounded variation, and the available extension theorems cannot be used†.

The problems we face vary, depending on the context. There are 'ultraviolet' problems springing from the infinite divisibility of the paths or 'histories' we are trying to sum over, and there are 'infrared' problems arising in connection with histories that are unbounded in time. By limiting ourselves to spatio-temporally discrete processes, we nullify the former problems, and that will be the context of the rest of this paper, where we will only encounter discrete histories like those that occur in a random walk. We will thus only be occupied by issues of infinite time.

The concrete instances we will consider will be of one of two types, which we can characterise by the kind of 'sample space' or 'history space', Ω, that we build on. The first instance arises in the context of quantum gravity and, more specifically, within the causal set programme. There the discreteness reflects the finiteness of Planck's constant, and the underlying physical process is a kind of 'birth' or 'accretion' process by means of which the causal set is built up or 'grows'. The corresponding sample-space of 'completed' causal sets consists of all the countable, past-finite partial orders P; and we seek to define a certain type of vector-valued measure μ on it.‡ In the second type of example, the elements of Ω will be discrete-time trajectories moving in a lattice that will be either the integers modulo n ($\mathbb{Z}_n$) or just the integers as such ($\mathbb{Z}$). These examples correspond to a widely studied class of processes known as 'quantal random walks', but for us they will be important primarily as simplified analogues of causal set growth processes. In that role, they are particularly illuminating because their sample spaces are essentially the same as those investigated in 'descriptive set theory'.

But how certain are we that the quantal measures in all these instances really need to be defined in a new way? With the lattices, the dynamical laws in question are those of the evolution generated by a unitary operator or 'transfer matrix'. In their

† See unpublished notes *Path Integrals* by R. Geroch available at `http://www.perimeterinstitute.ca/personal/rsorkin/lecture.notes/geroch.ps`.)

‡ The dynamics determines μ only up to a unitary transformation. The object of direct physical interest is not μ itself but a certain scalar-valued set-function belonging to the class of strongly positive decoherence functionals or quantal measures on Ω. However, any such functional can be represented (Dowker *et al.* 2010) as a measure on Ω that is valued in some Hilbert space $\mathfrak{H}$.

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

path-integral formulation, such unitary laws inevitably lead to measures of unbounded variation (Dowker *et al.* 2010), and Kolmogorov–Carathéodory-type theorems are thus guaranteed to fail. In the more important, causal set case, however, there remains some doubt, especially given the anticipated breakdown of unitary evolution in that case. The only fully developed dynamics we have for causal sets is that of the classical sequential growth (CSG) models, which in themselves are not quantal in nature. For them, the usual extension theorems do suffice because we are dealing with a classical probability measure (Brightwell *et al.* 2003). But if we complexify the parameters of a CSG model, we straightforwardly obtain a family of quantal measures (decoherence functionals) that are in general neither unitary nor of bounded variation (Dowker *et al.* 2010). Although none of these complexified CSG dynamics is likely to exhibit quite the type of interference required by quantum gravity, the fact that the measures that arise are not of bounded variation suggests that this might turn out to be a general feature of quantal causal sets, just as it is a general feature of quantal path integrals in other contexts. Nevertheless, it is worth bearing in mind the possibility that the physically appropriate quantal measures for causal set dynamics will turn out to be σ-additive in the traditional sense. Were that to happen, quantum gravity would have revealed itself to be more tractable mathematically than the nominally much 'simpler' non-relativistic free particle! The problems addressed in the present paper would then be pseudo-problems, as far as quantum gravity went.

In our current state of ignorance, however, it seems prudent not to count on so much good fortune. And, besides, we might still like to have a well-defined path-integral for systems like the free particle, but without having to embed them in a full-blown theory of quantum gravity. What, then, can we do when bounded variation fails? As argued in Dowker *et al.* (2010), such a failure need not be the end of the story because in a concrete physical situation, the space of histories has more structure than is available in an arbitrary measure space. Indeed, physicists routinely work with infinite sums and integrals that converge only conditionally. Typically, they introduce a 'cutoff' or integrating factor in a manner mandated by physical considerations, in effect, doing the sums or integrals in a particular order so that their convergence need not be absolute. In the case of the planar disk, for example, instead of expressing it as a disorganised sum of an infinite number of rectangles, we might think to employ a definite sequence of approximations, each consisting only of rectangles bigger than a certain size ϵ. The area would then be given by the $\epsilon \to 0$ limit of these approximations.

In our situation, we can attempt something similar by considering 'late-time' cylinder sets to be 'finer' than 'early-time' ones. In order to implement this idea, we will first, for any given set $A \subseteq \Omega$ of histories or paths, look for a 'canonical' sequence of approximations A_n to A in terms of cylinder sets, and this sequence should be as near to unique as feasible. Then, given such a sequence, we will try to decide what further convergence properties it ought to have in order that we can form a limit $\mu(A)$ of the individual $\mu(A_n)$ and consistently attribute this limit to A as its quantal measure.

In the following, we will only take a few steps in the direction indicated, pointing out along the way various pitfalls we will need to avoid. Hopefully, this can at least illustrate the kind of approach we might take to the extension problem for quantal measures. Sections 6 and 7 are the heart of the paper. In Section 6, we will define canonical approximations for a limited class of *events*, as (sufficiently regular) subsets $A \subseteq \Omega$ are normally designated. We will then introduce a convergence criterion for an approximating sequence A_n in Section 7, and prove that the resulting extension of μ is additive for disjoint unions of open sets. Limited as our approximation scheme will be, it will at least embrace the type of event A that is most important for the sake of causal sets, namely the *covariant stem event*. Among all the sets of histories we might wish to assign a measure to, these are the only indispensable ones. Without them, it would be nearly impossible to produce a generally covariant dynamical scheme in any useful sense (Brightwell *et al.* 2002; Brightwell *et al.* 2003).

An appendix lists some of the symbols used in the body of the paper.

2. Sample-spaces and amplitudes for causal sets and the 2-site hopper

2.1. *Causal sets*

A causal set (or *causet*) (Bombelli *et al.* 1987; Sorkin 2005; Henson 2009; Dowker 2006; Surya 2011) in its most general form can be any locally finite partial order or *poset*, but in the context of the dynamics of sequential growth and quantal cosmology, no element of the causet will possess more than a finite number of ancestors. For our present purposes, we may thus define a causet as a *past-finite countable poset*, that is, a countable (possibly finite) set of *elements* endowed with a transitive, acyclic order-relation, $\prec$, which I will also take to be irreflexive. These concepts are described in greater detail in Rideout and Sorkin (2000), where the notion of sequential growth is also explained. Here we will just summarise the main definitions and introduce the notation we will use.

A sequential growth process proceeds as a succession of 'births' of new elements, and in this sense is never ending. If, however, we idealise it as having 'run to completion', it will have produced a *completed causet* as defined above: a countable set of elements, each having a finite number of predecessors or *ancestors* but a possibly infinite number of descendants. The set of all such causets constitutes the natural sample-space Ω for this process. In fact, we must distinguish here two distinct sample spaces, which we may call Ω^{gauge} and $\Omega^{physical}$. The latter, which in some sense is the true sample space, consists of *unlabelled* causets, or, equivalently, isomorphism equivalence classes of causets. The former, which we will normally simply denote by Ω, then consists of the *naturally labelled* causets, where a natural labelling is a numbering, $0, 1, 2, \ldots$, of the elements that is compatible with the defining order $\prec$: if $x \prec y$, then y carries a bigger label than x. Here again, of course, we really intend isomorphism equivalence classes of labelled causets (or, if you like, the elements could be taken to be the integers themselves in this case).

The labels record the order of the respective births, and what is most important for us here is that this order is supposed to be fictitious in the same sense as a choice of coordinate system for a continuous space–time is fictitious. The physically meaningful or *covariant* events will thus correspond to subsets of $\Omega^{physical}$, whereas the measure μ defining the growth process is in the first instance defined on Ω^{gauge}. But even a very simple subset of $\Omega^{physical}$, even a singleton, will equate to a much less accessible subset of Ω^{gauge}, namely, the subset obtained by taking every possible natural labelling of every

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

member of the original subset. In this way, the need arises for an extension of μ that will assign well-defined measures to such 'covariant' subsets of Ω^{gauge}. Unlike the case for the example of the hopper to be discussed next, this is not just a matter of convenience if we want to be in a position to ask truly label-independent questions about the causet.

For the rest of this paper, all causets will be labelled unless otherwise specified. (Brightwell *et al.* (2003) used Ω to denote the true or 'covariant' sample space, and $\widetilde{\Omega}$ for its labelled counterpart). Here, however, it seems simpler to use Ω for the latter since it is the space we will usually be dealing with.

In Brightwell *et al.* (2003), the measures defining the CSG dynamical models were defined rigorously by extending a probability measure given originally on the space $\mathfrak{Z}$ of *cylinder events* (or cylinder sets), where a cylinder event $\text{cyl}(c) \in \mathfrak{Z}$ is by definition the set of all completed causets containing a given, naturally labelled, finite causet c. A finite causet will also be called a *stem*, and, on occasion, a 'truncated history'. In conjunction with these definitions, we will also define $\Omega(n)$, which is the space of all naturally labelled causets of n elements, and $\mathfrak{Z}(n)$ or $\mathfrak{Z}_n$, which is the space of cylinder events of the form $\text{cyl}(c)$ for $c \in \Omega(n)$. The cylinder sets comprise what is called a 'semiring' of sets in the sense that given any two cylinder sets, Z_1 and Z_2, their intersection, $Z_1Z_2 \equiv Z_1 \cap Z_2$, is also a cylinder event, and their difference $Z_1 \backslash Z_2$ is the disjoint union of a finite number of cylinder events. In fact, the cylinder events form an especially simple kind of semiring because any two of them are either disjoint or nested.

To go through the definition of the CSG models in general would take us too far afield, but the special case of 'complex percolation' is simple enough to be given here as an illustration of the general scheme. The vector measure μ is determined in this case by a single complex parameter p, and it takes its values in a one-dimensional Hilbert space, which we may identify with $\mathbb{C}$, so that $\mu(A)$ is itself just a complex number. Now let $c \in \Omega(n)$ be a labelled causet of n elements and let $Z = \text{cyl}(c)$ be the corresponding cylinder set. Then $\mu(Z) = p^L(1-p)^I$, where $L = L(c)$ is the number of links in c and $I = I(c)$ is the number of incomparabilities. Here an incomparability is simply a pair of unrelated elements, and a link is a causal relation, $x \prec y$, which is 'nearest neighbour' in the sense that there exists no intervening z for which $x \prec z \prec y$.

Observe now that the collection of naturally labelled finite causets, that is, the space $\bigcup_n \Omega(n)$, has itself the structure of a poset in a natural way. Indeed this poset is actually a tree $\mathfrak{T}$, because its elements are labelled. (The corresponding structure formed by the unlabelled stems is a more interesting poset called *poscau* in Rideout and Sorkin (2000).) Clearly, a particular realisation of the growth process, or, equivalently, the resulting completed causet in Ω, can be thought of as an upward path through this tree. An analogous concept will be possible for the two site hopper, and in this guise seems to have been heavily exploited in descriptive set theory (Kechris 1995; Moschovakis 2009). (See figures 1 and 2.)

Finally, we will define an *event algebra* to be a family of subsets of the sample space Ω closed under the operations of intersection and complementation. An event algebra is thus a Boolean algebra or 'ring of sets'. To the extent that it can be achieved, we normally want the domain of μ to be such an algebra, because, for example, if the events 'A happens' and 'B happens' are of interest, then so also is the event 'either A or B happens'.

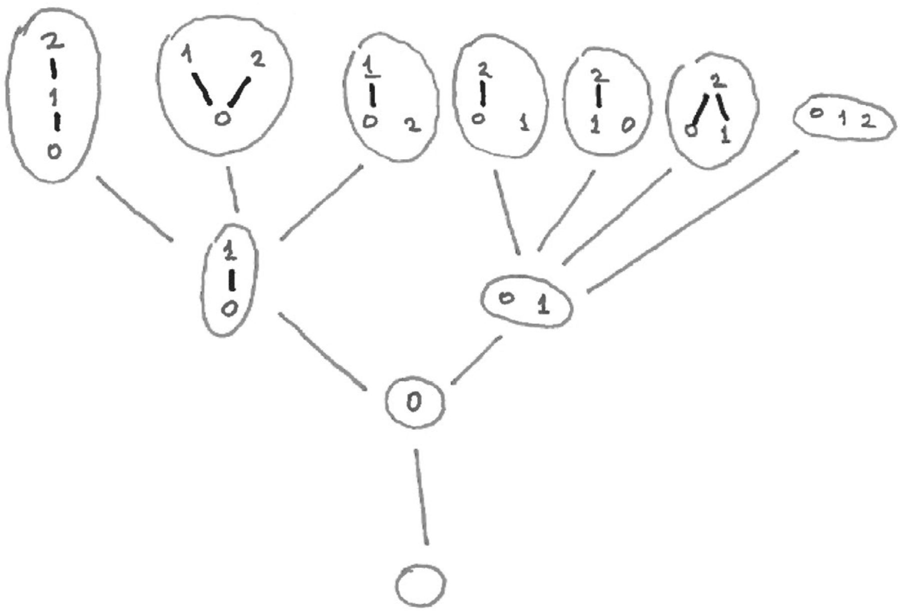

Fig. 1. The first 4 levels of the tree $\mathfrak{T}$ of naturally labelled causets

The cylinder sets $\mathfrak{Z}$ do not themselves form an algebra, but the family $\mathfrak{S}$ of finite unions of cylinder sets does – in fact, it is $\mathfrak{R}(\mathfrak{Z})$ the Boolean algebra generated by $\mathfrak{Z}$. In all cases of interest, μ will automatically extend uniquely from $\mathfrak{Z}$ to $\mathfrak{S}$, yielding a finitely additive measure on it. The space $\mathfrak{S}$ thus constitutes a minimum domain of definition for the vector measure μ. The question then will be how far μ can be extended beyond $\mathfrak{S}$ into the σ-algebra generated by it, the hope being that the enlarged domain $\mathfrak{A}$ will itself be an event algebra, and that it will contain enough events such that, at a minimum, the physically most important questions will become well posed. (Some noteworthy instances of covariant questions/events will be discussed in the next section.)

2.2. *n-site hopper*

By '2-site hopper' we mean the formalisation of a particle residing on a 2-site lattice and at each of a discrete succession of moments, it either stays where it is or jumps to the other site (Gudder and Sorkin 2012). For definiteness, we will assume that the moments are labelled by the natural numbers and the sites by $\mathbb{Z}_2$, and that the hopper begins at site 0 at moment 0. The definitions of sample space, cylinder event, and so on are closely analogous to those given above for causets, and references to them should be understandable for the moment without their formal definitions, which will be given after the transition amplitudes have been specified. The full course of the motion, which is idealised as having run to completion, will be called a *path* or 'history'. Notice that, modulo the small ambiguity in how a real number can be expressed as a 'binary decimal', each such path can be identified uniquely with a point in the unit interval $[0, 1] \subseteq \mathbb{R}$.

Aside from having a simpler sample space than in the causet case, the hopper also offers us a fuller illustration of the problems of defining the vector measure corresponding to a path integral. Unlike the former case, where the correct choice of quantal amplitudes is only conjectural, there exists for the hopper a choice that can be interpreted as a straightforward discretisation of the Schrödinger dynamics of a non-relativistic free particle moving on a circle (*cf.* Pearle (1973)).

These amplitudes can be understood more easily if we set them up, not just for two sites, but for the more general case of the circular lattice $\mathbb{Z}_n$ ('n-site hopper'). They may look more familiar if we present them as the unitary evolution operator or 'transfer matrix' analogous to the propagator that solves the Schrödinger equation in the continuous case. To that end, let $x \in \mathbb{Z}_n$ be the location of the particle at some moment t, let x' be its location at the next moment $t' = t+1$, and for brevity write $\exp(2\pi i z) \equiv \mathbf{1}^z$. The amplitude to go from x to x' in a single step is then

$$\frac{1}{\sqrt{n}}\,\mathbf{1}^{(x-x')^2/n}$$

for n odd, and

$$\frac{1}{\sqrt{n}}\,\mathbf{1}^{(x-x')^2/2n}$$

for n even. For example, for $n = 6$ and with $q = \mathbf{1}^{1/12}$, the (un-normalised) amplitudes to hop by 0, 1, 2 or 3 sites are $q^0 = 1$, $q^1 = q$, q^4, and $q^9 = -i$, respectively. For the 2- and 3-site hoppers, the above amplitudes are particularly simple, yielding for $n = 3$ the transfer matrix

$$\frac{1}{\sqrt{3}}\begin{pmatrix} 1 & \omega & \omega \\ \omega & 1 & \omega \\ \omega & \omega & 1 \end{pmatrix} \qquad\qquad (\omega = 1^{1/3})$$

and for $n = 2$ the transfer matrix

$$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix}. \tag{1}$$

From these expressions and the definition of the decoherence functional, it is not hard to construct the equivalent vector measure along the lines of Dowker *et al.* (2010). In the simplest case of two sites, which will be our main example here, μ is valued in a two-dimensional Hilbert space $\mathbb{C}^2$ and, with a convenient choice of basis vectors, can be expressed as follows. Let $(0\,x_1\,x_2\,x_3 \ldots x_m)$ be a truncated path and $Z \subseteq \Omega$ be the corresponding cylinder event. Then $\mu(Z) \equiv |Z\rangle$ will be the two-component complex vector v_α where (with no summation implied)†

$$v_\alpha = (U^{-m})_{\alpha x_m} U_{x_m x_{m-1}} \cdots U_{x_3 x_2} U_{x_2 x_1} U_{x_1 0}, \tag{2}$$

U being the unitary matrix of equation (1). Note, incidentally, that U^j is periodic with period 8 and is very easy to compute explicitly, since $U^4 = -1$ while $U^2 = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}$ is also very simple.

† Another notation for v_α could be $\langle \alpha | 0\,x_1\,x_2\,x_3 \ldots x_m \rangle$

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

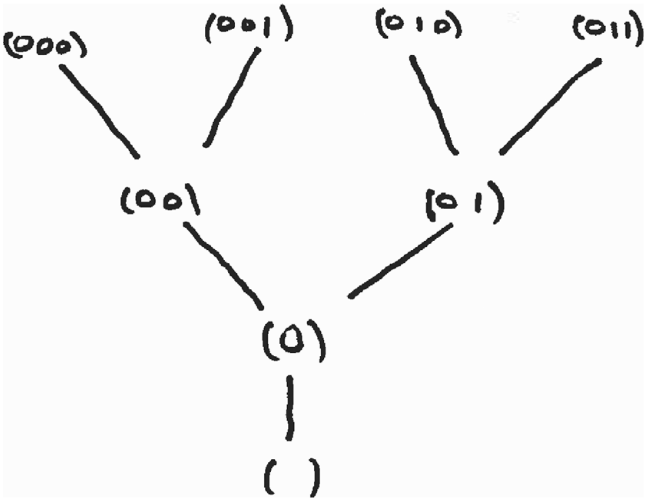

Fig. 2. The first 4 levels of the tree $\mathfrak{T}$ for the 2-site hopper

Finally, we will give the formal definitions for the 2-site hopper. A truncated history, the counterpart of a finite causal set, is for the hopper an initial segment of a path, for example, $(0, 1, 1, 0, 1)$. (Recall our boundary condition that all paths begin at zero.) The set of all such truncated histories with length n when the initial 0 is omitted will be $\Omega(n)$, and the corresponding cylinder events will be the elements of $\mathfrak{Z}_n$. The semiring $\mathfrak{Z}$ will be the union of the $\mathfrak{Z}_n$. For example, $\text{cyl}(0, 1, 1, 0, 1) \in \mathfrak{Z}_4$ is the set of all completed paths of the form $(0, 1, 1, 0, 1, x_5, x_6, \cdots)$. Exactly as above, $\mathfrak{S}$ will be the Boolean algebra generated by $\mathfrak{Z}$. We can check straightforwardly that μ, as defined by (2), extends uniquely and consistently to each $\mathfrak{S}_n$, and therefore to $\mathfrak{S}$ as a whole. Again, the truncated histories can be construed as the nodes of a tree $\mathfrak{T}$, the 'branches' or 'edges' being given by path extension. (See Figure 2.) For example, there will be an edge from $(0, 1, 1, 0)$ to $(0, 1, 1, 0, 1)$.

In the following, it will sometimes be enlightening to consider hopper-paths on the infinite lattice $\mathbb{Z}$. In that case, the paths will be restricted to move no more than one site per step ('random walk') so that the resulting tree $\mathfrak{T}$ continues to have a finite number of branches emanating from each node.

For an extensive discussion of the quantal 2-site hopper, see Gudder and Sorkin (2012). For more general sorts of quantal random walks, see Martin *et al.* (2005).

3. Some events whose measures we would like to define

The event algebra $\mathfrak{S}$ generated by the cylinder events supplies enough events to allow us to ask any question† about the process under consideration, as long as it does not refer to happenings arbitrarily far into the future. But it is often the case that we do not want

† The words 'event' and 'question' are in a certain sense synonyms. For an event $A \subseteq \Omega$ there is a corresponding question 'Does A happen?'. Note in this connection that (except in the classical case) it would lead to confusion if we read 'A happens' as 'the path is an element of A' – *cf.* Sorkin (2012).

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

to be bound by this limitation, especially since in the causet case the 'time' referred to contains a large element of gauge, as explained above.

To illustrate how 'infinite-time' events enter the story, we will now look at a few examples, beginning with the n-site hopper. Perhaps the simplest and most familiar example of this kind is the event R of *return*, which occurs if and when the particle returns to its starting point at some later time. This event, in other words, is the set of all paths $(0\,x_1\,x_2\,\ldots)$ for which one of the $x_i = 0$. Plainly, R is not in $\mathfrak{S}$, because the return, although it must occur at a finite time if it occurs at all, can take place at an arbitrarily late time. For a *classical* hopper on a finite lattice, we know that $\mu(R)$, the measure of the return event (which classically is its probability), is unity, but in order to express this fact directly, we need R to be in the domain of μ. Of course, we could avoid any direct reference to R by introducing the finite-time event R_n that the particle returns on or before the nth step. Instead of asserting that $\mu(R) = 1$, we could then say 'the sequence $\mu(R_n)$ converges to 1 as $n \to \infty$'. Plainly, the first formulation is simpler and less cumbersome to work with. Notice in this case that not only at the level of the measures, but even at the level of the events themselves, R is the limit of the R_n in a natural sense, since the latter are *nested* and 'increase monotonically to R'. That is, we have $R_1 \subseteq R_2 \subseteq R_3 \cdots$, with R itself being the union of the R_n, or, logically speaking, their 'disjunction'. Were μ a classical measure, this would guarantee convergence of the $\mu(R_n)$ and consistent extension of the domain of μ to include the event R; in the quantal case it guarantees nothing.

A similar event to 'return', but one that is related even less directly to any cylinder event, is the event R^∞ that the particle visits $x = 0$ infinitely often. This event also has a well-defined probability of unity in the classical case. Since, however, it cannot come to fruition at any finite time, it cannot, unlike the event R of simple return, be expressed as a union of cylinder sets or other members of $\mathfrak{S}$. Instead, it is a countable intersection of events, each of which is a countable union of events in $\mathfrak{S}$. For example, let $E(j,k)$ for $j < k$ be the event that $x_k = 0$. Then $R^\infty = \bigcap_j \bigcup_k E(j,k)$. (In words, for each moment j there is a later moment k at which the particle visits the origin[†].) To give meaning to $\mu(R^\infty)$ by prolonging the initially defined measure with domain $\mathfrak{S}$, we would have to think in terms of a limit of limits.

As a third example (restricted this time to one of the lattices, $\mathbb{Z}$ or $\mathbb{Z}_n$ with $n > 4$), consider the event that the particle visits $x = 3$ but never reaches $x = 5$. This is intermediate between the two previous examples in its remoteness from $\mathfrak{S}$, and is naturally expressed as the set-theoretic difference of two limits of finite-time events, the first being, naturally, the event F that the particle reaches $x = 3$, and the second, G being the event that it reaches $x = 5$. Just as with the return event R, the event $F \backslash G$ is, in a well-defined sense to which we will return below, a limit of events in $\mathfrak{S}$, but it is not simply the union or intersection of a monotonically increasing or decreasing sequence.

[†] We can often arrive at such combinations by beginning with a formal statement of what it means for the event to happen. In this case, we might first write down what it means for R^∞ *not* to occur: $(\exists n_0)(\forall n > n_0)(x_n \neq 0)$, and then negate it to obtain $(\forall n_0)(\exists n > n_0)(x_n = 0)$. The nested combination of unions and intersections is basically just a translation of this second statement into set-theoretic language.

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

It is useful at this point to introduce some further notation to help us discuss the types of events we have just met. Let $\mathfrak{X}$ be any collection of subsets of Ω closed under pairwise union and intersection. Then $\bigvee \mathfrak{X}$ will be the family of events of the form $\bigcup_{n=1}^{\infty} X_n$, where $X_n \in \mathfrak{X}$ and $X_1 \subseteq X_2 \subseteq X_3 \cdots$. It is easy to see that $\bigvee \mathfrak{X}$ is also closed under union and intersection, and also that it would not change if we dropped the monotonicity condition, $X_1 \subseteq X_2 \subseteq X_3 \cdots$. In words, the members of $\bigvee \mathfrak{X}$ are the unions of monotonically increasing events in $\mathfrak{X}$. For the intersections of monotonically decreasing events in $\mathfrak{X}$, we will write dually $\bigwedge \mathfrak{X}$. And for the Boolean algebra generated by $\mathfrak{X}$, we will write, as above, $\mathfrak{R}\,\mathfrak{X}$ or $\mathfrak{R}(\mathfrak{X})$. Our first example, 'return', is then an element of $\bigvee \mathfrak{S}$, our second of $\bigwedge \bigvee \mathfrak{S}$, and our third of $\mathfrak{R} \bigvee \mathfrak{S}$, while for $\mathfrak{S}$ itself, we have $\mathfrak{S} = \mathfrak{R}(\mathfrak{Z})$.

Turning now to events for causal sets, we will encounter some types that are very similar to those just discussed. Foremost in importance are the unlabelled stem events mentioned earlier. Given two causets c and c', of which the first is finite, we say that c' admits c as a *stem* (or 'partial† stem') if c' contains a downward-closed subset that is isomorphic to c. In the context of sequential growth, this can also be expressed by saying that it might have happened that elements of c were all born before any of the remaining elements of c'. A stem thus generalises the notion of an 'initial segment'. The *stem event* 'stem(c)' is then the set of all $c' \in \Omega$ that admit c as a stem. The stem c that enters this definition is taken to be unlabelled because our aim is to produce a label-independent or 'covariant' event. It is evident that stem(c) is indeed covariant in this sense, since the condition that defines it does not refer to the labelling of c'.

The importance of the stem events physically is that, essentially, any covariant question that we care to ask about the causet can in principle be phrased in terms of stem events. The precise result proved in Brightwell *et al.* (2003) is that every covariant event is equal, up to a set of measure zero, to a member of the σ-algebra generated by the stem events. We can also prove that any covariant event that is open in the topology of Section 4 later in this paper is a countable union of stem events, which is a purely topological result that holds independently of any assumption about the measure μ. Ideally then, the domain of μ would embrace the whole σ-algebra generated by the stem events. At a minimum, one would hope that it would embrace the stem events themselves.

Now the event 'stem(c)' does not belong to the domain $\mathfrak{S}$ on which μ is initially defined because it is not a finite-time event when referred to 'label-time'. If it were, there would exist some integer N such that if the growing causet c' admitted c as a stem, it would already admit it as soon as the first N elements had been born. But, in fact, there is nothing in principle to stop the stem in question appearing at an arbitrarily late stage of the growth process. Evidently, the situation is like that of the hopper event 'return'. Based on this analogy, we would expect the stem events to be found in $\bigvee \mathfrak{S}$, and so they are, and this follows directly from the fact that any stem event is a union of cylinder sets:

$$\text{stem}(c) = \bigcup \left\{ \text{cyl}(\widetilde{b}) \in \mathfrak{Z} \,|\, \widetilde{b} \text{ admits } c \text{ as a stem} \right\}. \tag{3}$$

† We can also define 'full stems' (Rideout and Sorkin 2000), but there is no special reason to consider them here.

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

The problem of extending the vector measure μ from $\mathfrak{S}$ to $\bigvee \mathfrak{S}$ is thus the most basic one for causal sets.

Starting from the stem events, we can build up other covariant events, whose occurrence or non-occurrence is of interest for cosmology. The simplest of these is the event that the causet is 'originary', meaning that all its elements descend from a unique minimal element or 'origin'. To say that a completed causet is originary is simply to say that it contains no second minimal element, for which it is necessary and sufficient that it fails to admit the 2-element antichain as a stem. (An antichain is a set of elements that are mutually unrelated or 'spacelike' to one another.) Thus, the event 'originary' is the complement of the event 'stem(a)', where a is the antichain of two elements. As such, it belongs to $\bigwedge \mathfrak{S}$, since, as we have seen, the stem events all belong to $\bigvee \mathfrak{S}$, and union turns into intersection under complementation.

If an originary causet represents a certain kind of 'big bang', then a causal set containing what is called in the combinatorics literature a *post* describes a 'cosmic bounce'. (A post is an element of a poset that is spacelike to no element.) In its degree of remoteness from the elementary cylinder events, the post event is comparable to the event of 'infinite return' in the case of a random walk, the similarity being even closer if we compare the post event to the *complement* of the infinite return event. In fact, both events belong to $\bigvee \bigwedge \mathfrak{S}$, although it is less easy to demonstrate this for the post event than it is in the case of return. To see why this is, nevertheless, true, imagine watching a succession of births of causet elements, $x_0, x_1, x_2 \ldots$ and waiting for a post to be born. If the birth in question is that of element x_n, then x_n must have every previous element as an ancestor: $x_j \prec x_n$ for all $j < n$. This renders x_n momentarily a 'candidate for becoming a post', but it does not guarantee that x_n will remain a viable candidate forever. In order for that to occur, every subsequent element, $x_{n+1}, x_{n+2}, \cdots$, must arise as a descendant of x_n, that is, $x_j \succ x_n$ for all $j > n$. By thinking of the post event P in this way, namely, as the set of all sequences of births satisfying the condition that a candidate post appear at some stage n and then not lose its viability at any later stage $m > n$, we can deduce that P belongs to $\bigvee \bigwedge \mathfrak{S}$. The most 'covariant' (albeit not the most direct) construction along these lines proceeds by first expressing P in terms of stem events; this will also illustrate the thesis that all covariant questions of interest can be expressed in terms of stem events.

Proceeding in this way, note first that if x is a post, then its exclusive past $T = \{y \mid y \prec x\}$ is not only a stem, but what has been called a 'turtle' (Varadarajan and Rideout 2006), meaning in the present context a stem that wholly precedes its complement: $(\forall x \in T)(\forall y \notin T)(x \prec y)$. Some thought reveals that a causet contains a turtle of n elements if and only if every stem of cardinality $n + 1$ has a unique maximal element. Introducing the term *principal* for such a stem, together with the terms n-stem (respectively, n-turtle) for a stem (turtle) of n elements, we can say succinctly that *a causet contains an n-turtle if and only if every $(n + 1)$-stem is principal.* Furthermore, it is easy to demonstrate that x is a post if and only if both its exclusive *and* inclusive pasts are turtles (the exclusive past being $\{y \neq x \mid y \prec x\}$ and the inclusive past being $\{y \mid y \preceq x\}$). Therefore, in a labelled causet, element x_n is a post if and only if every stem of either $n + 1$ or $n + 2$ elements is principal. Let P_n be the event that this happens, then the post event itself is $P = \cup_{n=1}^{\infty} P_n$.

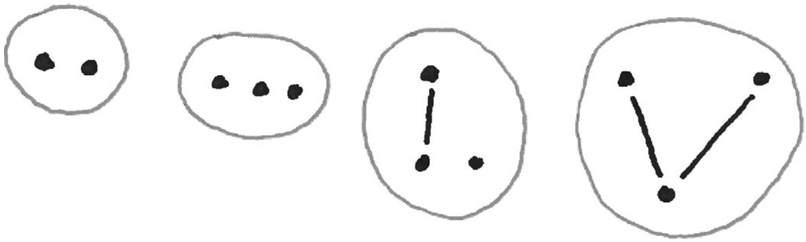

Fig. 3. The non-principal 2- and 3-stems.

Now let us examine the event P_n more closely. It *fails* to happen if and only if some $(n+1)$-stem or $(n+2)$-stem fails to be principal. Let $S_1^n, S_2^n, \ldots, S_{K_n}^n$ be an enumeration of all such stems (there being only a finite number of n-stems, for any n), and let $Q_j^n = \text{stem}(S_j^n)$ be the corresponding stem events. We then obtain P_n in the 'manifestly covariant' form $P_n = \Omega \backslash (\cup_j Q_j^n) = \Omega \backslash (\cup_j \text{ stem}(S_j^n))$. P is thus a countable union of finite Boolean combinations of stem events:

$$P = \bigcup_{n=0}^{\infty} (\Omega \setminus (\bigcup_{j=1}^{K_n} \text{stem}(S_j^n))). \tag{4}$$

If we knew how to take stem events as primitive, P would thus be a rather simple type of event, inasmuch as the inner union only ranges over a finite number of events. But given that the existing dynamical schemes all begin with labelled causets, we will still need to trace everything back to the cylinder events $\mathfrak{Z}$.

First, though, a simple example might be in order, say for $n = 1$. (The event P_0 is just the originary event, which we might not even want to count as a post.) The event P_1 requires that all 2- and 3-stems be principal. The only 2-stem that can occur is thus the 2-chain $(a \prec b)$, while the admissible 3-stems are the 3-chain $(a \prec b \prec c)$ and the 'Λ-order' $(a \prec c, b \prec c)$. The stems that must be excluded – those denoted above by S_j^n – are correspondingly the 2- and 3-stems that are not principal: the 2-antichain, the 3-antichain, the 'L-order' $(a \prec b, c)$ and the 'V-order' $(a \prec b, a \prec c)$ – see figure 3.

To complete the demonstration that $P \in \bigvee \bigwedge \mathfrak{S}$, let us return exclusively to labelled causets, observing first that, in view of equation (4), it suffices to show that the complement of a finite union of stem events belongs to $\bigwedge \mathfrak{S}$.† To this end, recall that any stem event A is an increasing union of events in $\mathfrak{S}$. Its complement, $\Omega \backslash A$, is therefore a decreasing intersection of complements of events in $\mathfrak{S}$, each of which is itself in $\mathfrak{S}$ since the latter, being a Boolean algebra, is closed under complementation. Hence, the complement of a stem event belongs to $\bigwedge \mathfrak{S}$, and the same holds for the complement of a finite union of stem events, such as occurs in (4).

† Strictly speaking, given the way we have defined the operation $\bigvee$, we also need to convert the outer union in (4) into an *increasing* countable union of events in $\bigwedge \mathfrak{S}$. That this is possible follows readily from the relation (6) of Section 5, which informs us that, when A_n and B_n are both decreasing sequences of sets, the union of their limits coincides with the limit of the decreasing sequence $A_n \cup B_n$, in consequence of which, $\bigwedge \mathfrak{S}$ is closed under finite union, and we can replace a countable union $\cup_n F_n$ of events $F_n \in \bigwedge \mathfrak{S}$ with the increasing union $\cup_n F'_n$, where $F'_n = \cup_{m \leqslant n} F_n$.

For reasons that will become clear shortly, it is natural to designate the elements of $\mathfrak{S}$ as *clopen*, meaning 'both closed and open' in the sense of point-set topology†. The events in $\bigvee \mathfrak{S}$ will then be *open*, those of $\bigwedge \mathfrak{S}$ will be *closed*, and we will have expressed our post event $P \subseteq \Omega$ as an increasing limit of closed subsets of Ω. Continuing in this vein, more elaborate combinations of the clopen events can be formed, including, for example, the event that infinitely many posts occur. But the physical relevance of such combinations seems to shrink rapidly as their complexity grows. Indeed, we might feel that, questions of convenience aside, no event more complicated than a finite Boolean combination of stem events can claim to be indispensable. We might even go further and call into doubt the status of complementation (negation), leaving unquestioned only those events formed as finite unions and intersections of stem events.

4. Ω as a compact metric space

4.1. *Open and closed sets*

Whenever the idea of convergence plays a role, we can expect, almost by definition, that topology will make an appearance. In the present situation we are talking about convergence to a given event $A \subseteq \Omega$ of a sequence of approximating events A_n, where in the first instance the A_n are formed as finite unions of cylinder events and thus belong to the event algebra $\mathfrak{S}$. In setting up such a sequence of approximations, we would like, as explained earlier, to regard those cylinder events that specify a greater portion of the history as more 'fine grained' than those that specify a lesser portion. This leads very naturally to a definition of distance between histories that makes Ω into a compact metric space (Brightwell *et al.* 2003; Kechris 1995).

In the case of causal set growth processes, the definition runs as follows. For each pair of completed labelled causets $a, b \in \Omega$, we set

$$d(a, b) = 1/2^n, \tag{5}$$

where n is the largest integer for which the elements $a_0\, a_1 \cdots a_n$ produce the same poset (with the same labelling) as elements $b_0\, b_1 \cdots b_n$. It is easy to verify that this yields a metric on Ω. Indeed, d satisfies a condition stronger than the triangle inequality: for any three causets a, b and c, we have $d(a, c) = \max(d(a, b), d(b, c))$. This 'ultrametric' property follows from the tree structure of the space $\mathfrak{Z}$ of cylinder sets (or equivalently truncated histories), as described earlier. The maximum distance between two causets is $1/2$, and occurs when their initial two elements already form distinct partial orders. Notice also that the open balls in this metric are exactly the cylinder sets, with the radius of the ball serving as a measure of 'fineness'.

It is not difficult to see that with this metric, Ω becomes a compact topological space†. Moreover, the cylinder sets, being the balls of some radius, are both open and closed, that

† In the context of abstract measure theory, the term 'elementary sets' was used in Kolmogorov and Fomin (1961) to refer to events analogous to those of $\mathfrak{S}$.

† This will be proved explicitly in the next section.

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

is, clopen. It follows by definition that Ω has a basis of clopen sets and that every open set is a countable union of cylinder sets (there being only a countable number of cylinder sets because the finite causets are only countable in number). Since each element of the event algebra $\mathfrak{S}$ is itself a (finite) union of cylinder sets, we can conclude that the open sets are precisely the members of $\bigvee \mathfrak{S}$, the closed sets, their complements, then being the members of $\bigwedge \mathfrak{S}$. The events that belong to both these families are the clopen events, and they clearly include all of $\mathfrak{S}$, because a finite union of open (respectively, closed) sets is also open (respectively, closed). We will prove the converse, that is, that every clopen event belongs to $\mathfrak{S}$.

Lemma 4.1. $\bigvee \mathfrak{S} \cap \bigwedge \mathfrak{S} = \mathfrak{S}$

Proof. We are asked to prove that $\mathfrak{S}$ comprises precisely the clopen subsets of Ω. Since we already know that every $A \in \mathfrak{S}$ is clopen, it suffices to verify that any clopen A, that is, any $A \in \bigvee \mathfrak{S} \cap \bigwedge \mathfrak{S}$, also belongs to $\mathfrak{S}$. Let $A \in \bigvee \mathfrak{S}$. By definition, A is a union of cylinder sets:

$$A = Z_1 \cup Z_2 \cup Z_3 \ldots.$$

Now, this sequence either terminates at a finite stage or it does not. If it terminates, then A is a finite union of cylinder sets, and thus a member of $\mathfrak{S}$, and we are done. If it does not terminate, we can find a sequence of points $x_j \in A$ that escape from every Z_j, and because Ω is compact, we can suppose that this sequence converges to some $x \in \Omega$. This x cannot lie in any given Z_k because it is a limit of points x_j that eventually belong to the closed set $\Omega \backslash Z_k$. Consequently, $x \notin A = \cup_k Z_k$. We have thus constructed a sequence of points of A that converge to a point outside A, meaning that A is not closed, and thus cannot be clopen. In other words, if it is clopen, we are back to the terminating sequence and the conclusion that $A \in \mathfrak{S}$. □

Turning to the 2-site hopper, we only need to make one change to what we wrote above for causets. The histories are now sequences of digits, 0 or 1, beginning with 0, and the integer n that occurs in the definition (5) is now the largest index such that the two subsequences $(0\ a_1\ a_2 \cdots a_n)$ and $(0\ b_1\ b_2 \cdots b_n)$ coincide. The rest is all the same. The history space Ω is still a compact metric space, the cylinder sets are clopen and generate the topology, and so on.

4.2. *The tree of truncated histories*

We have already seen in Section 2, that a point of Ω, that is to say a history, can be construed as a path γ through the tree $\mathfrak{T}$, each node of which is a 'truncated history', meaning, depending on the context, either a finite causet or a finite sequence of binary digits†. From this correspondence between histories and paths through $\mathfrak{T}$, we get an

† The paths under consideration in the following will usually begin at the 'root' of $\mathfrak{T}$ (corresponding to the cylinder set Ω), but sometimes they will start at some other node of $\mathfrak{T}$. The two cases are actually interchangeable because any path not starting at the root has a unique extension back to it since $\mathfrak{T}$ is a tree.

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

alternative way to characterise certain types of events, including the open sets $\bigvee \mathfrak{S}$ and, more generally, the events in $\mathfrak{R} \bigvee \mathfrak{S}$.

We first consider a cylinder set $Z = \text{cyl}(h)$, where $h \in \Omega(n)$ is a truncated history and ask which paths γ correspond to this cylinder set? By definition, they are just the paths whose corresponding histories reproduce h when truncated at the nth stage, that is, they are precisely the paths that *pass through the node* in $\mathfrak{T}$ that represents h, which we will either denote by h itself or by $\text{node}(h)$ in order to emphasise the fact that h is being treated as a node in $\mathfrak{T}$. Because $\mathfrak{T}$ is a tree, such a path necessarily follows one of the branches emanating from $\text{node}(h)$, and then remains for ever in the 'upward subset' of $\mathfrak{T}$ consisting of all descendants (in $\mathfrak{T}$) of h. In this way, every open event $A \subseteq \Omega$ can be represented by an upward-closed subset $\alpha \subseteq \mathfrak{T}$, and *vice versa*, given such a subset the paths that enter (and consequently remain in) α comprise an open event $A \subseteq \Omega$. More generally, every subset $\alpha \subseteq \mathfrak{T}$ gives rise to an event $S(\alpha)$ by the same rule.

Definition 4.2. $S(\alpha) = \{\gamma \mid \gamma \text{ is eventually in } \alpha\}$.

Here, γ is a point of Ω, represented as a path $\gamma = (h_0, h_1, h_2 \cdots)$ through $\mathfrak{T}$, and the statement that this path is *eventually* in α means that $(\exists n_0)(\forall n > n_0)(h_n \in \alpha)$. It is clear that $S(\cdot)$ commutes with the Boolean operations:

$$S(\alpha\beta) = S(\alpha)S(\beta)$$
$$S(\alpha \backslash \beta) = S(\alpha) \backslash S(\beta)$$
$$S(\alpha + \beta) = S(\alpha) + S(\beta),$$

and so on (where $\alpha + \beta := (\alpha \cup \beta) \backslash (\alpha\beta)$ is the Boolean operation of 'addition modulo 2').

As a further aid to intuition, we can think of certain types of events in terms of 'properties' acquired or lost in the course of the process under consideration. Formally, this corresponds closely to the characterisation by sets of nodes in $\mathfrak{T}$, but it has a more 'evolutionary' feel to it. For example, consider the event of return analysed earlier. We can cook up a 'property' that the particle possesses when, and only when, it has returned to the origin. By definition, this property of 'having returned' is *hereditary* in the sense that, once acquired, it can never be lost. Topologically, the set of all paths γ that acquire a hereditary property yields an open subset of Ω – it is easy to corroborate this by thinking through the definitions. Dually, a property that if it is once lost can never be regained, but that every path begins with, corresponds to a closed set (a causet example of this is being originary). And a property that can be acquired but never regained if lost yields an event of the form $A \backslash B$, where both A and B are open (a hopper example of this is visiting $x = 1$ but not $x = 2$). Notice that this third type of property includes both of the previous two as special cases. In terms of sets of nodes like the sets α discussed above, the first type of property is an upward-closed subset of $\mathfrak{T}$, the second is a downward-closed subset and the third is a *convex* subset, defined as a subset of $\mathfrak{T}$ that contains, together with nodes h_1 and h_2, every node that lies on some path from h_1 to h_2. In order-theoretic language for the poset $\mathfrak{T}$, this just says that α includes the *order-interval* between any two

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

of its elements.[†] In Sections 5 and 6, the events of the form $S(\alpha)$ for some convex α will be among those for which we will able to produce a canonical representation as a limit of clopen events.

As an application of some of these ideas, we can prove the assertion made earlier that Ω is topologically compact. By a standard criterion for compactness, it suffices to prove that any covering of Ω by cylinder sets has a finite sub-covering, so consider an arbitrary collection of cylinder sets $Z \in \mathfrak{Z}$ that covers Ω. In relation to $\mathfrak{T}$, such a covering is a collection of nodes that no path γ can avoid forever. Now the (incomplete) paths that do avoid these nodes fill out a subtree[‡], $\mathfrak{T}'$ of $\mathfrak{T}$, with the property that no path γ can remain within $\mathfrak{T}'$ forever. But it is well known that such a tree can only have a finite number of nodes, assuming that no node has an infinite branching number (this has been called the 'infinity lemma' of graph theory). The maximal elements of $\mathfrak{T}'$ thus furnish a finite collection of nodes that every path must encounter. In their guise as cylinder sets, these nodes constitute a finite subcover of Ω.

5. Set-theoretic limits of events

The most elementary kind of limit we can imagine for a sequence of events is not metric, topological or measure-theoretic, but purely set-theoretic. We have already seen how increasing sequences of clopen sets yield the open sets $\bigvee \mathfrak{S}$, while decreasing sequences of clopen sets yield the closed sets $\bigwedge \mathfrak{S}$. In both cases the relevant limit concept emerges more or less automatically. Going beyond these two types of approximation, we can recognise a more general concept, of which $\bigvee$ and $\bigwedge$ are special cases. We let X_j be a sequence of subsets of Ω, and say that it is convergent if we have for any point x of Ω that eventually $x \in X_j$ or eventually $x \notin X_j$. In such a case, we will write $X = \lim X_j$, where, of course, X consists of those x that realise the first alternative of being eventually in X_j. We will write $\mathrm{Lim}\,\mathfrak{S}$ for the set of all events obtainable in this way as limits of events $A_j \in \mathfrak{S}$. Notice that 'lim' commutes with the Boolean operations:

$$\lim(A_n \cup B_n) = (\lim A_n) \cup (\lim B_n), \quad \text{and so on.} \tag{6}$$

In trying to extend our vector measure μ beyond the clopen events, we might hope that we could at least get as far as $\mathrm{Lim}\,\mathfrak{S}$. Were μ an ordinary measure, this would be true, because $\lim A_j$ would be sandwiched between the measurable sets

$$\limsup A_j = \cap_j \cup_{k>j} A_k$$
$$\liminf A_j = \cup_j \cap_{k>j} A_k,$$

both of which are equal to $\lim A_j$ when the latter exists. This would ensure that $\lim A_j$ was measurable and that

$$\lim \mu(A_j) = \mu(\lim A_j).$$

† The order-interval delimited by elements x and y of some poset is $\{z \mid x \prec z \prec y\}$.
‡ That is, a downward-closed subset of $\mathfrak{T}$.

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

But this argument is not available with quantal measures, and it turns out that convergence can fail already for certain decreasing sequences of clopen events whose measures diverge to infinity (Gudder and Sorkin 2012). On the other hand, convergence succeeds for many other sequences, and we might hope that the failures were confined to physically uninteresting questions.

Of course, the failure of convergence in even some cases is likely to contaminate other cases, which makes us doubt whether $\mu(A)$ can be defined without some further limitation on the sequence A_j beyond the mere requirement that $\lim A_j = A$. In the next two sections we will investigate some restrictions of this sort. For the moment, we will just note that for open sets A there exists a very naturally defined canonical sequence of events $A_n \in \mathfrak{S}_n$ converging to A. Namely, we can take for A_n the union of all the cylinder sets from $\mathfrak{Z}_n$ that are contained within A. This yields a 'best approximation to A at stage n' in the sense that A_n could not be enlarged without the sequence losing its increasing nature.

Dually, we immediately obtain a canonical choice of sequence for any closed event B (just apply complementation to the sequence of clopen events approximating $\Omega \backslash B$). However, having two different classes of canonical sequences in this way introduces an ambiguity for events that are both open and closed. Fortunately, the ambiguity in this case does no harm because a clopen event necessarily belongs to $\mathfrak{S}$ according to Lemma 4.1, so both the increasing and decreasing canonical sequences terminate at a finite stage: they differ only transiently.

Leaving aside questions of convergence and uniqueness, we might wonder how many events the above limit process can access, even in the best case. That is, how many of the interesting questions even belong to $\operatorname{Lim} \mathfrak{S}$? With reference to the causal set case, we first recall that we encounter all the stem events without ever leaving the open sets $\bigvee \mathfrak{S}$. Remembering also that $\operatorname{Lim} \mathfrak{S}$ is closed under the Boolean operations, we can thus say on the positive side that every finite logical combination of stem events is available within $\operatorname{Lim} \mathfrak{S}$ (as also is the entire event algebra $\mathfrak{R} \bigvee \mathfrak{S}$ of course). On the negative side, however, events like the post event and (for the particle case) the event of infinite return fall outside $\operatorname{Lim} \mathfrak{S}$ as a consequence of the following lemma.

Lemma 5.1. Let $A \subseteq \Omega$. If both A and $\Omega \backslash A$ are dense subsets of Ω, then $A \notin \operatorname{Lim} \mathfrak{S}$.

Proof. In the following, we write $A \perp B$ to mean that A and B are disjoint. Suppose, in order to show a contradiction, that $A = \lim A_n$ with $A_n \in \mathfrak{S}$, and write $\overline{A}_n$ for its complement $\Omega \backslash A_n$, also taking note of the fact that $\overline{A}_n$, like A_n itself, is clopen. We will find inductively a subsequence $A_{n_1}, A_{n_2}, A_{n_3}, \cdots$ of the A_n and a matched sequence of clopen sets $B_1 \supseteq B_2 \supseteq B_3 \cdots$ such that B_j is alternately included in and disjoint from A_{n_j}:

(1) We begin by putting $n_1 = 1$ and $B_1 = A_{n_1}$, so we have $B_1 \subseteq A_1$.
(2) Since B_1 is open and $\Omega \backslash A$ is dense, there exists $x \in B_1 \cap (\Omega \backslash A)$. Then, since $x \notin A = \lim_n A_n$, there exists by hypothesis some $n_2 > n_1$ such that $x \notin A_{n_2}$, that is, $x \in \overline{A}_{n_2}$. We now put $B_2 = \overline{A}_{n_2} \cap B_1$, which is again clopen since both $\overline{A}_{n_2}$ and B_1 are clopen, so $B_2 \subseteq B_1$ with $B_2 \perp A_{n_2}$.
(3) We now proceed exactly as in Step (2), but with the roles of A and $\Omega \backslash A$ interchanged. Specifically, since B_2 is open and A is dense, there exists $x \in B_2 \cap A$. Then, since

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

$x \in A = \lim_n A_n$, there exists by hypothesis some $n_3 > n_2$ such that $x \in A_{n_3}$. We now put $B_3 = A_{n_3} \cap B_2$, which is again clopen since both A_{n_3} and B_2 are clopen, so $B_3 \subseteq B_2 \subseteq B_1$ with $B_3 \subseteq A_{n_3}$.

We now proceed inductively to produce $B_4 \subseteq A_{n_4}$, $B_5 \perp A_{n_5}$, and so on. Finally, we put

$$B = \lim_n B_n = \bigcap_{n=1}^{\infty} B_n$$

and note that B is non-empty since the B_n are all compact (indeed, every event in $\mathfrak{S}$ is compact, being a closed subset of the compact space Ω). We now pick any $x \in B$. For odd j, we have

$$x \in B_j \subseteq A_{n_j} \Rightarrow x \in A_{n_j},$$

and for even j, we have

$$x \in B_j \perp A_{n_j} \Rightarrow x \notin A_{n_j}.$$

Thus the A_n alternate between including and excluding x, contradicting our assumption that $\lim A_n$ exists. □

The lemma applies to the post event because no matter how far the growth process has proceeded, the growing causet 'still has a free choice' of whether to end up with or without a post (and exactly the same thing can be said for the event of infinite return). But this freedom means precisely that both the post event and its complement are dense in Ω.

Lemma 5.1 shows that $\operatorname{Lim} \mathfrak{S}$ is a long way from containing every event of potential interest, but we might wonder exactly how far. One answer comes from Kechris (1995, Exercise 22.17), according to which, $\operatorname{Lim} \mathfrak{S}$ equals what is called Δ^0_2, which is defined to be the intersection of $\bigvee \bigwedge \mathfrak{S}$ and $\bigwedge \bigvee \mathfrak{S}$. This places $\operatorname{Lim} \mathfrak{S}$ at a very low level of the so-called 'Borel hierarchy', which continues on for $\aleph_1$ steps beyond Δ^0_2 before it exhausts the Borel subsets of Ω. In this sense the limiting process 'lim' does not take us very far beyond the clopen events. On the other hand, we have also seen that by applying 'lim' more than once, we can reach, for example, the post event. How many events can we reach in this manner? Exercise 22.17, in combination with other results in Kechris (1995), also answers this question by implying that (transfinite but still countable) iteration of the 'lim' operation suffices to produce any Borel set. In this sense, the lim operation is quite far reaching, given that it would be hard to conceive of an event of interest that does not fall within the Borel domain.

6. Canonical approximations for certain events

We have already discovered one canonical sequence A_n of approximations to A for an event $A \subseteq \Omega$ that is open with respect to the topology defined in Section 4, that is, for $A \in \bigvee \mathfrak{S}$. The cylinder sets $\mathfrak{Z}_n$ 'at stage n' provide a kind of 'mesh' in Ω whose fineness increases with n, and our canonical choice of approximating event at stage n was

$$A_n = \bigcup \{ Z \in \mathfrak{Z}_n \,|\, Z \subseteq A \}, \tag{7}$$

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

which is the biggest member of $\mathfrak{S}_n = \mathfrak{R}\mathfrak{Z}_n$ that can fit inside A. As we have seen, the A_n converge to A in the sense defined in section 5, but, of course, there are many other sequences $B_n \in \mathfrak{S}_n$ that also converge to A in this sense, and when the vector measure μ is not of bounded variation, there is no guarantee that the corresponding sequences $\mu(A_n)$ and $\mu(B_n)$, if they converge at all, will converge to the same limit. In general, they doubtless will not if B_n is chosen with sufficient malice. In the face of such ambiguity, we might still hope to find some reasonably inclusive event algebra $\mathfrak{A} \supseteq \mathfrak{S}$, and for each event $A \in \mathfrak{A}$ a canonical approximating sequence of events $A_n \in \mathfrak{S}_n$ with $\lim A_n = A$ such that $\mu(A_n)$ is a convergent sequence in Hilbert space. The vector $\lim_n \mu(A_n)$ could then be adopted as the definition of $\mu(A)$.

One snag that this perspective encounters is already apparent for the case where we are approximating open sets A and B, and our canonical approximations A_n and B_n are the ones given by equation (7). From $\lim A_n = A$ and $\lim B_n = B$, it does indeed follow, as we have already noted, that

$$\lim(A_n \cap B_n) = A \cap B$$
$$\lim(A_n \cup B_n) = A \cup B.$$

For the case of intersection, it even follows that the events $(A_n \cap B_n)$ provide the canonical approximations to the open event $A \cap B$, but the analogous conclusion fails for the case of union because the canonical approximation $(A \cup B)_n$ will, in general, be larger than $(A_n \cup B_n)$, since some cylinder set $Z \in \mathfrak{Z}_n$ can, by 'straddling the boundary' between A and B, be included in $A \cup B$ without being included in either A or B. We would thus obtain different approximating sequences for $A \cup B$ depending on whether we regard it as an open set in its own right or as the result of uniting A with B. In the next section we will begin to see what it would take to render this kind of ambiguity harmless. For now, however, we will ignore that issue and simply consider the question of finding unambiguous approximating sequences for as many members of $\mathfrak{R} \bigvee \mathfrak{S}$ $(= \mathfrak{R} \bigwedge \mathfrak{S})$ as possible.

To that end, we will return to the tree $\mathfrak{T}$ of truncated histories and the method of representing certain events by subsets $\alpha \subseteq \mathfrak{T}$. Although we did not make it explicit earlier, it is clear that a sequence of events $A_n \in \mathfrak{S}_n$ is equivalent to a set of nodes $\alpha \subseteq \mathfrak{T}$. Indeed, each A_n is a union of cylinder sets $Z \in \mathfrak{Z}_n$, and each such cylinder set corresponds to a node in $\mathfrak{T}_n$, the nth level of $\mathfrak{T}$. This associates with each A_n a set of nodes at level n, and amalgamating the nodes of all levels into a single collection yields α. Conversely, given $\alpha \subseteq \mathfrak{T}$, we obtain A_n as the union of the cylinder sets that correspond to the α-nodes at level n. Since the correspondences between cylinder sets $Z \in \mathfrak{Z}_n$, nodes in $\mathfrak{T}_n$ and truncated histories $\gamma \in \Omega(n)$ are so close, we will often identify all three with one another, and speak, for example, of a cylinder set Z as a node in $\mathfrak{T}$. When this is done, we can express the correspondence between node-sets α and approximating sequences (A_n) using a simple formula by writing $A_n = \bigcup(\alpha \cap \mathfrak{Z}_n)$.

Now let α be any set of nodes and let the A_n be the corresponding sequence of events. Recall that we defined $S(\alpha)$ as the event that γ is *eventually* in α:

$$S(\alpha) = \{\gamma \mid (\exists N)(\forall n > N)(\gamma_n \in \alpha)\}.$$

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

Dually, we can also define $\widetilde{S}(\alpha)$ as the event that γ is *repeatedly* in α:

$$\widetilde{S}(\alpha) = \{\gamma \mid (\forall N)(\exists n > N)(\gamma_n \in \alpha)\}.$$

It then follows straightforwardly from the definitions that

$$\begin{aligned} S(\alpha) &= \liminf A_n \\ \widetilde{S}(\alpha) &= \limsup A_n. \end{aligned} \tag{8}$$

Since $\lim A_n$ exists if and only if $\liminf A_n = \limsup A_n$ (in which case their common value equals $\lim A_n$), we learn that the events of the form $\lim A_n$ are precisely those for which $S(\alpha) = \widetilde{S}(\alpha)$, which in turn are precisely those such that no path γ can leave and re-enter α more than a finite number of times. It is clear that this property generalises the concept of convexity we encountered earlier. Notice, incidentally, that equations (8) imply that the forms $S(\alpha)$ and $\widetilde{S}(\alpha)$ do not reach beyond $\bigvee\bigwedge\mathfrak{S}$ and $\bigwedge\bigvee\mathfrak{S}$, which are known in descriptive set theory as Σ^0_2 and Π^0_2, respectively. Roughly speaking, they reach as far as events whose complexity is that of the post event. Very optimistically, we might hope to go beyond this and find for any Borel set $A \subseteq \Omega$ some sort of canonical presentation in terms of clopen events, but in this section we will not venture outside of $\mathfrak{R}\bigvee\mathfrak{S}$, the finite Boolean combinations of opens. Since $\mathfrak{R}\bigvee\mathfrak{S} \subseteq \mathrm{Lim}\,\mathfrak{S}$, all such events can be expressed as $S(\alpha)$ for some subset $\alpha \subseteq \mathfrak{T}$.

What we are asking for is a sort of 'normal form' for events E in $\mathfrak{R}\bigvee\mathfrak{S}$. As a first step in that direction, we will prove that every such event can be expressed as a disjoint union of events, each of which has the form open$\backslash$open, or, equivalently, open $\cap$ closed.

Lemma 6.1. Let $E \in \mathfrak{R}\bigvee\mathfrak{S}$ be a finite logical combination of open events. Then there exists a decreasing sequence of open events $E^1 \supseteq E^2 \supseteq E^3 \cdots \supseteq E^K$ such that

$$E = E^1 + E^2 + E^3 \cdots + E^K = E^1\backslash E^2 \sqcup E^3\backslash E^4 \sqcup \cdots,$$

where '$\sqcup$' denotes disjoint union. Moreover, the E^j are formed from the original events using only the operations of union and intersection.

Proof. In the proof, as in the statement of the lemma, we use the operation of Boolean addition,

$$A + B = (A \cup B)\backslash(A \cap B), \tag{9}$$

and write the intersection of two sets as their product. Any Boolean combination of sets is then a polynomial in these sets, and since products of open sets are open, any Boolean combination of open events can be expressed simply as a Boolean sum of open events. Given these facts, it is not hard to devise a proof by induction, but we will just illustrate the pattern involved using the cases $K = 2, 3$. For two events, we have

$$\begin{aligned} A + B &= (A + B + AB) + AB \\ &= A \cup B + AB, \end{aligned}$$

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

and for three, we have

$$\begin{aligned} A+B+C &= A+(B+C) \\ &= A+(B\cup C+BC) \\ &= (A+B\cup C)+BC \\ &= (A\cup B\cup C+A(B\cup C))+BC \\ &= A\cup B\cup C+(A(B\cup C)+BC) \\ &= A\cup B\cup C+A(B\cup C)\cup BC+A(B\cup C)BC \\ &= A\cup B\cup C+(AB\cup AC\cup BC)+ABC. \end{aligned}$$

The 'inclusion–exclusion' pattern evident here emerges with particular clarity when we interpret Boolean addition as the addition of characteristic functions modulo 2. The final equation in the statement of the lemma then follows directly from the fact that the E^j are decreasing. We can also restate the essence of the proof in a simple formula:

$$\sum_{j=1}^{K} A_j = \sum_{j=1}^{K} B_j,$$

where

$$B_j = \{x \,|\, x \text{ belongs to at least } j \text{ of the } A_k\},$$

which is clearly a union of intersections of the A_k. □

Given any set E expressed as in the lemma, we immediately get the approximations

$$E_n = E_n^1 + E_n^2 + \cdots + E_n^K,$$

where E_n^j is our canonical nth approximation to the open set E^j, and thence the corresponding sets of nodes

$$\alpha_n = \alpha_n^1 + \alpha_n^2 + \cdots + \alpha_n^K,$$

together with their union $\alpha = \cup_n \alpha_n$. However, this construction is only a first step toward uniqueness because the resulting α still depends on the original choice of the E^j, which are not given to us uniquely by the lemma.

In working toward a unique approximating sequence, we will concentrate on the simplest case of an event $E = A\backslash B$, which is the difference of only two open sets $B \subseteq A$ (corresponding to $K = 2$ in the lemma). We are then led to ask if we can render A and B unique in this case. It is not difficult to demonstrate that if we gather together all pairs $A \supseteq B$ such that $E = A\backslash B$, then the union of all the sets A and the union of all the sets B yields another such pair. Evidently this 'biggest pair' is unique and uniquely determined by the original event E. This in turn yields, by (6), a canonical sequence of approximations E_n to E of the form, $E_n = A_n\backslash B_n$, where $A_n \in \mathfrak{S}_n$ and $B_n \in \mathfrak{S}_n$ are the canonical nth approximations to A and B. In terms of the equivalent node-sets α_n, these approximations are given by $\alpha_n\backslash\beta_n$, whose union over n we will simply denote by α, following our earlier notation.

Although the node-set α we have found is canonical and concisely defined, we might hope to find a more constructive route to it, or at least a characterisation of it in terms of more easily verifiable necessary and sufficient conditions. The rest of this section will develop a prescription of this sort. In fact, I am not certain that the second prescription will be strictly equivalent to the first. If it is, that is all to the good since we will then not be forced to choose between the two. Otherwise, it does not really matter since the second prescription stands on its own and, being more concrete, is likely to be more useful in practice.

Lemma 6.2. If α and β are upward-closed subsets of $\mathfrak{T}$ with $\alpha \supseteq \beta$, then $\alpha\backslash\beta$ is convex.

Proof. We need to show that no path between two nodes x and y in $\alpha\backslash\beta$ can contain nodes outside of $\alpha\backslash\beta$. Equivalently, no path that has left $\alpha\backslash\beta$ can ever re-enter it. But since α is upward-closed no path from $x \in \alpha$ can leave α, therefore it can leave $\alpha\backslash\beta$ only by entering β, and it then must remain in β (which is also upward-closed) forever, and, consequently, can never re-enter $\alpha\backslash\beta$. □

Now let $E = A\backslash B$ be as above, with A and B open, and let α and β be the corresponding node-sets. Since A and B are open, both α and β are upward-closed subsets of $\mathfrak{T}$. We also know that $A = S(\alpha)$, $B = S(\beta)$ and $A\backslash B = S(\alpha\backslash\beta)$. The lemma then tells us that $A\backslash B = S(\widehat{\alpha})$, with $\widehat{\alpha}$ a convex subset of $\mathfrak{T}$. The converse is also true, as shown by the following lemma.

Lemma 6.3. If $\widehat{\alpha} \subseteq \mathfrak{T}$ is convex, then $S(\widehat{\alpha}) = A\backslash B$ for some open sets A and B.

Proof. Recall that we have identified points of Ω with infinite paths γ through $\mathfrak{T}$, and let A be the set of all paths that enter $\widehat{\alpha}$, and let B be the subset of these that subsequently leave $\widehat{\alpha}$. By definition, $S(\widehat{\alpha}) = A\backslash B$, but both A and B are open because the property of 'having entered $\widehat{\alpha}$' and the property of 'having left $\widehat{\alpha}$' are both hereditary. □

From now on, we will just deal with the convex subset $\widehat{\alpha}$, and rename it to plain α for simplicity. That is, we will be concerned with a fixed event E of the form (open\open) and with a convex set of nodes $\alpha \subseteq \mathfrak{T}$ such that† $E = S(\alpha)$.

We will say that $\alpha \subseteq \mathfrak{T}$ is *prolific* if it lacks maximal elements. A second very natural requirement in addition to convexity is the condition that α be prolific in this sense. Given convexity, this is equivalent to saying that every node $x \in \alpha$ originates a path that remains forever within α. In the opposite case, α will contain 'sterile' nodes from which all paths eventually leave α for good. It is clear that removing these sterile nodes will not alter E, nor will it spoil the convexity of α. We can therefore always arrange that α be both convex and prolific. The 'pruning' of the 'sterile' nodes to render α prolific also appears as a very natural operation when it is expressed in terms of cylinder sets Z. It simply removes from α those Z that are disjoint from E.

We have now arranged for α to be convex and prolific, but this does not yet make it unique, since, for example, we could remove all the nodes up to any fixed finite level n without altering $S(\alpha)$. If we did so, however, we might create a situation where, for

† In view of (8), we would, in general, also want to require that $S(\alpha) = \widetilde{S}(\alpha)$, but this holds automatically when α is convex.

example, some cylinder set Z was wholly included in E without Z itself (regarded as a node in $\mathfrak{T}$) belonging to α. To remedy this sort of lacuna, we can adjoin to α every node Z such that every path originating at Z eventually enters α. It is again easy to see that adjoining these nodes will not interfere with α being convex and prolific.

In this last step we have, in a manner of speaking, completed α toward the past, but, in fact, there is cause to carry this process of 'past-completion' farther by adjoining yet more nodes to α. These additional nodes are perhaps not such obvious candidates as the previous ones, but throwing them in as well (which I think corresponds to enlarging the open set A) will provide us with the uniqueness we are seeking.

Definition 6.4. Let $x \in \mathfrak{T}$ and $\alpha \subseteq \mathfrak{T}$. Then $x \prec \alpha$ means that x precedes some node in α:

$$(\exists y \in \alpha)(x \prec y).$$

Remark 6.5. In terms of cylinder sets, $Z_1 \prec Z_2 \iff Z_2 \supseteq Z_1$.

Definition 6.6. The *exclusive past* of α is the set of nodes strictly below α: $\{x \notin \alpha \mid x \prec \alpha\}$.

Using this definition, we will say that α is *past-complete* if its exclusive past P is prolific, which in turn says that any node in P originates a path that repeatedly visits P. I claim we can render α past-complete be adjoining to it all nodes below α that fail to satisfy this last condition, and, furthermore, that the resulting set of nodes α' will yield the same event E as α and will be convex and prolific if α itself was.

Lemma 6.7. Let $\alpha \subseteq \mathfrak{T}$ be any set of nodes and α' be its 'past-completion' as just described. Then α' is past-complete. Moreover, $S(\alpha') = S(\alpha)$ and $\widetilde{S}(\alpha') = \widetilde{S}(\alpha)$.

Proof. The fact that $S(\alpha') \supseteq S(\alpha)$ is obvious. To prove that they are equal, it suffices to show that no path can eventually remain within α' without also remaining eventually within α. To show a contradiction, we suppose the contrary to be true, and let $\gamma \in S(\alpha')$ be a path that is repeatedly outside α. By passing to a tail of γ, we can suppose that it is always within α'. Let $x \in \gamma$ be a node that is not in α and let $y \in \gamma$ be a later node of the same type. Since $y \in \alpha' \backslash \alpha$, it is by definition in P, the exclusive past of α. Hence, x originates a path (namely γ) that visits P at y; and since there are an infinite number of nodes like y, γ visits P repeatedly. But this contradicts the criterion for having included x in α' in the first place.

The proof that $\widetilde{S}(\alpha') = \widetilde{S}(\alpha)$ is similar. It suffices to show that every path in $\widetilde{S}(\alpha')$ visits α repeatedly. To show a contradiction, we suppose the contrary to be true, and let $\gamma \in \widetilde{S}(\alpha')$ be a path that is eventually outside α. By passing to a tail, we can suppose that γ is always outside of α. Let $x \in \gamma$ be a node that is in α' and let $y_1, y_2, \cdots$ be a sequence of later nodes of γ that are also in α'. Since the y_j belong to $\alpha' \backslash \alpha$ they are by definition in P, the exclusive past of α. Hence x originates a path that returns repeatedly to P, which contradicts the criterion for having included x in α' in the first place.

To complete the proof, we need to show that α' is past-complete. To that end, let x be in the exclusive past of α'. Since any y in α' is either in α or in its exclusive past, and since $\alpha' \supseteq \alpha$, we have x is also in the exclusive past of α. Consequently, since x was not put into α', it originates a path γ that repeatedly visits the exclusive past of α. But by definition, no node in such a path would have been put into α' either, so γ repeatedly visits the exclusive past of α', as we wished to show. □

We also wish to prove that past-completion preserves the attributes of being prolific and convex. The first part is easy because past-completion only adds nodes that are below some element of α, and this can introduce no new maximal element.

For the second part, we need to demonstrate[†] that if α is convex, and if $x \prec y$ are nodes in α' then the order-interval I delimited by x and y is also within α'. When $x \in \alpha$ the proof is simple, since y is either within α itself or precedes some element z that is. In either case, the interval I is included in some second interval I' (which is possibly the same as I) with endpoints in α. Then $I' \subseteq \alpha$ because α is convex, and thus

$$I \subseteq I' \subseteq \alpha \subseteq \alpha'$$

also, as desired. The remaining possibility is that $x \in \alpha' \backslash \alpha$, in which case it seems more convenient to deal with paths rather than intervals. From the definition of convexity, proving that α' is convex amounts to showing that no path originating from x can leave α' and then re-enter it. First, observe that since every node of α' precedes some node of α, no node of $\alpha' \backslash \alpha$ can follow a node of α, since if it did, it would also lie within the convex set α. In consequence, any path that exits α permanently, also exits α'. Now let γ be any path originating from $x \in \alpha' \backslash \alpha$, and, as before, write P for the exclusive past of α. By the definition of α', every path from x must eventually leave P. If it does so by leaving the past of α, $\{x \in \mathfrak{T} \,|\, x \prec \alpha\}$, then it certainly can never re-enter α'. If it does so by entering α, then it can exit α' only by exiting α, in which case it can never re-enter α or (as we have just observed) α'.

So far, we have established the existence, for our event E, of a node-set α that is convex, prolific and past-complete. We will now complete the story by proving that α is also unique. To that end, let α and β be two prolific, convex and past-complete node-sets such that $S(\alpha) = S(\beta)$. Does it follow that $\alpha = \beta$? In demonstrating that the answer is 'yes', we will use the *ad hoc* notation $\mathbb{P}(\alpha)$ for the exclusive past of α as defined earlier, with

$$\overline{\mathbb{P}}(\alpha) = \alpha \cup \mathbb{P}(\alpha) = \{x \in \mathfrak{T} \,|\, (\exists y \in \alpha)(x \preceq y)\}$$

being the *inclusive past*.

We will first establish that α and β have equal inclusive pasts: $\overline{\mathbb{P}}\alpha = \overline{\mathbb{P}}\beta$. In fact, if $x \in \overline{\mathbb{P}}\alpha$, then x originates a path γ that visits α. Since α is prolific, this path can be arranged to visit α repeatedly, and since α is convex, such a path can never leave α. Hence, $\gamma \in S(\alpha)$, implying, in particular, that γ visits β, so $x \in \overline{\mathbb{P}}\beta$. The converse follows symmetrically.

We now suppose, in order to show a contradiction, that there exists $x \in \alpha \backslash \beta$. Such an x belongs by definition to $\overline{\mathbb{P}}\alpha$, and hence to $\overline{\mathbb{P}}\beta$, and thus to $\mathbb{P}\beta$, which in turn is prolific by the definition of past-completeness. Thus x originates a path γ that repeatedly visits $\mathbb{P}\beta$.

[†] The demonstration that follows seems rather longer than it ought to be. Intuitively, it suffices to observe first that α' is built up from α by successive adjunction of maximal elements of its exclusive past, and then that adjoining such an element cannot spoil convexity.

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

We now claim that γ must leave α at some stage (since otherwise $\gamma \in S(\alpha) \Rightarrow \gamma \in S(\beta) \Rightarrow \gamma$ eventually in β, so γ could never again visit $\mathbb{P}\beta$). And since α is convex, γ must remain outside α once it has left. On the other hand, γ must continue to visit $\mathbb{P}\beta$, which in turn is a subset of $\overline{\mathbb{P}}\beta = \overline{\mathbb{P}}\alpha$. But if γ really visited some $y \in \overline{\mathbb{P}}\alpha$, then, by definition, we could divert it at y to some other γ' that would re-enter α, something that we have just proved to be impossible. This completes the proof of the following theorem.

Theorem 6.8. Every event E of the form $E = A \backslash B$ with A and B open can be expressed as $E = S(\alpha) = \widetilde{S}(\alpha)$ for a unique set of nodes α that is convex, prolific and past-complete.

The theorem furnishes a canonical set α of nodes corresponding to E, and as explained earlier, we immediately get from such an α a canonical sequence of approximants E_n to E such that $E = \lim_n E_n$. We have thus reached our immediate goal.

The canonical approximating sequences of the theorem provide a good reference point for further developments, and we have learned how to arrive at them step-by-step, starting from the open sets A and B. Nevertheless, it seems unlikely that we can limit ourselves to these particular approximants in general. Rather, as remarked already at the beginning of this section, we will, in general, have to deal with many different sequences converging to the same event, unless we can devise canonical sequences that are closed under the Boolean operations.

We have already encountered an ambiguity of this nature when we noticed that our original, increasing canonical approximants (7) for open events (we will call them 'C1') are not fully compatible with the Boolean operation of complementation. Specifically, for a clopen event E, these C1 approximants depend on whether we derive them directly from E or by complementing the corresponding approximants for the open event $\Omega \backslash E$. However, we run into a further, but related, conflict if we now compare the C1 approximants with those of the above theorem (we will call them 'C2'). For an open event E, the C1 approximant E_n is just the biggest member of $\mathfrak{S}_n$ included within E. But if we view E as the difference $E = A \backslash B$, with A being E itself and $B = 0$ being the empty event, the theorem provides a different set of approximants E_n. In general, the two disagree, as can be seen from the observation that the pair $(E\ 0)$ is not the 'biggest one' yielding E.

Consider, for example, the 2-site hopper event such that the particle does not remain forever at its starting site 0, but such that the first time it hops to site 1 it immediately returns to 0. This event is a union of cylinder sets corresponding to truncated trajectories of the shape $(0, 0, 0, \cdots 0, 1, 0, *)$ where the star '$*$' represents any finite sequence of zeros and ones. For this event, the node-set α_1 of type C1 consists of precisely the truncated trajectories just indicated. But that set of nodes is not past-complete. Its completion, the type C2 node-set α_2, also contains the truncated trajectories $(0, 0, 0, \cdots 0, 1)$. Notice that α_2 differs from α_1 at an infinite number of nodes in this case (Figure 4 illustrates this phenomenon).

We thus have to reckon with overlapping but, in general, incompatible prescriptions for different types of events. If one prescription were to be adopted exclusively, it should probably be C2, which covers more events than C1 does. (Incidentally, C2 resolves the aforementioned ambiguity in the C1 prescription in favour of treating clopen events as closed, not open.) We might make the counter-argument on behalf of C1 that monotonic

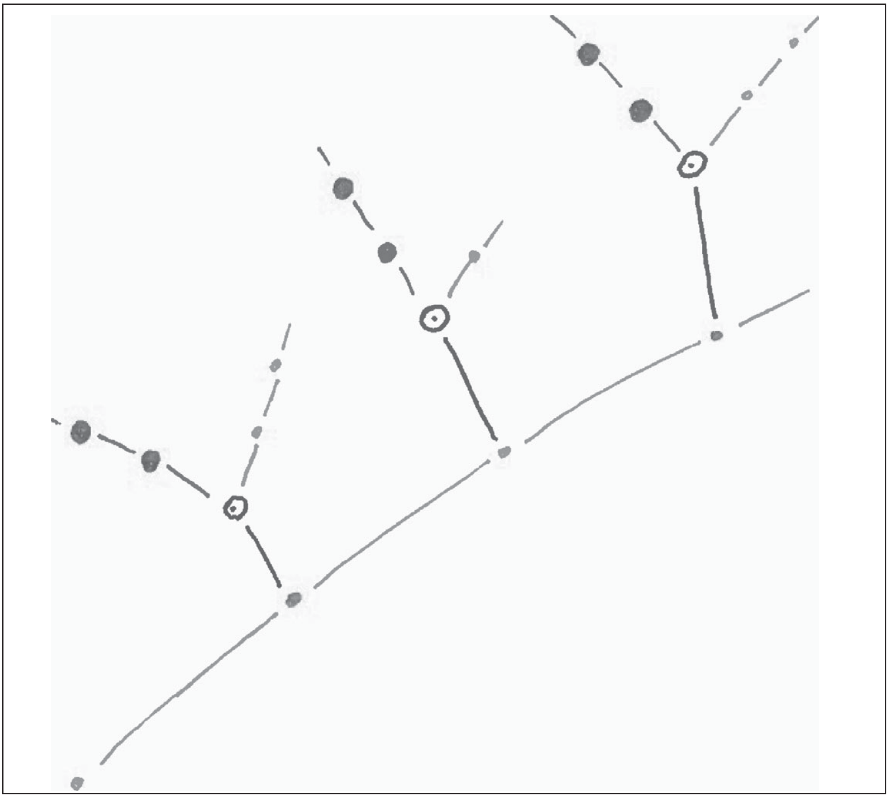

Fig. 4. Illustrating past-completion in the tree shown. The circled nodes 'complete' those above them shown as heavy dots. Let the node-set be α_1 before completion and α_2 after completion. It is clear that $S(\alpha_1) = S(\alpha_2)$ but only α_1 is upward-closed.

convergence of the approximants is to be preferred, but this does not seem so compelling in the context of a *quantal* measure, which itself is not a monotonic set-function. Better than either choice, however, would be not having to choose at all because the alternative approximating sequences would all lead to the same extension of our initial quantal measure. The main thing for now is that we have discovered at least one canonical choice of clopen events E_n converging to any event of the form $E = A \backslash B$ with A and B open.

In the face of these various ambiguities, we should emphasise that none of them affect, in the causet case, the stem events themselves, essentially because the latter are not only open but dense in Ω, or, more physically, because any growing causet that has not yet produced a given stem always retains a choice of whether or not to do so. It follows that not only does the C1 prescription coincide with the C2 prescription for stem events (its exclusive pasts being already prolific), but also the 'biggest pair' prescription with which we began provably agrees with the C1 prescription. The same ought to apply to finite unions and intersections of stem events, and similar comments could be made about the 'return' event in the hopper case.

We conclude this section by sketching very briefly how we might try to carry our successful 'canonisation' of $E = \text{open}\backslash\text{open}$ over to the general case where

$E = E^1 + E^2 \cdots + E^K$, the E^j being open and nested. Just as earlier we found a 'biggest pair', $A \supseteq B$, by forming unions of the individual events A and B, we can do the same thing here with the E^j to obtain (at least in principle) a canonical set of 'biggest' open and nested events E^j such that $E = E^1 + E^2 \cdots + E^K$. Associating its canonical approximants E^j_n (in the C1 sense, say) with each such E^j then yields for E itself the approximants $E_n = E^1_n + E^2_n \cdots + E^K_n$ such that $\lim E_n = E$. In principle, this achieves our goal, but it remains once again at a rather abstract level.

As before, we can attempt a more constructive development by working with the node-set $\alpha \subseteq \mathfrak{T}$ corresponding to our approximating sequence E_n, or perhaps with some similar node-set whose uniqueness can be established directly, and for which we can prove that $E = S(\alpha) = \widetilde{S}(\alpha)$. But how would our construction of α go in this more general case, and what would generalise the conditions that α be convex, prolific and past-complete? It is clear from Lemmas 6.2 and 6.3 that convexity is now too restrictive. In its place, we would probably try the more general requirement that no path γ could enter and leave α more than K times. Correspondingly, we might then expect that α would decompose into subsets α^j that were convex in the strict sense. We might also try to arrange for each α^j to be prolific and past-complete, hoping that this would again confer uniqueness on the whole collection. If all this worked out, we would have constructed a canonical approximating sequence for any Boolean combination of open events, and, in particular, for any Boolean combination of stem events.

A next step beyond $\mathfrak{R} \bigvee \mathfrak{S}$, if we could take it, would be to devise canonical approximations for larger families of events, starting with the collections $\bigvee \bigwedge \mathfrak{S}$ and $\bigwedge \bigvee \mathfrak{S}$ in which the post event and its complement are to be found. An event A in either of these collections is accessible from $\mathfrak{S}$ as a limit of limits, but such a double limiting process can only make the potential ambiguities worse. For example, the event R^∞ of repeated return is in $\bigwedge \bigvee \mathfrak{S}$. It could be expressed as $\lim A_n$, where $A_n =$ 'returns at least n times', or it could be expressed instead as $\lim B_n$, where $B_n =$ 'returns at least once after $t = n$'. Both A_n and B_n give decreasing sequences of open events and both converge to R^∞, but which sequence, if either, should be favoured as canonical? Perhaps, in certain cases, we could arrive at a canonical presentation by further generalising our treatment above in terms of sets of nodes $\alpha \subseteq \mathfrak{T}$, but beyond this, it is not easy to guess how we might proceed. To devise canonical approximations for events of still greater complexity would seem to demand a fresh approach.

Finally, it is worth repeating here that uniqueness in and of itself does not guarantee compatibility with the Boolean operations. And I believe, in fact, that none of the prescriptions that this section has considered are compatible with the full set of such connectives, though some are compatible, for example, with complement or disjoint union – *cf.* figure 5. If there did exist a compatible prescription, or even a prescription compatible 'modulo initial transients', which is just as good, it would weigh very heavily in its favour.

7. Evenly convergent sequences of events

We are given a vector-valued measure $\mu : \mathfrak{S} \to \mathfrak{H}$ defined initially on the finite unions of cylinder events, and we wish to enlarge this initial domain $\mathfrak{S}$ so that it can embrace events

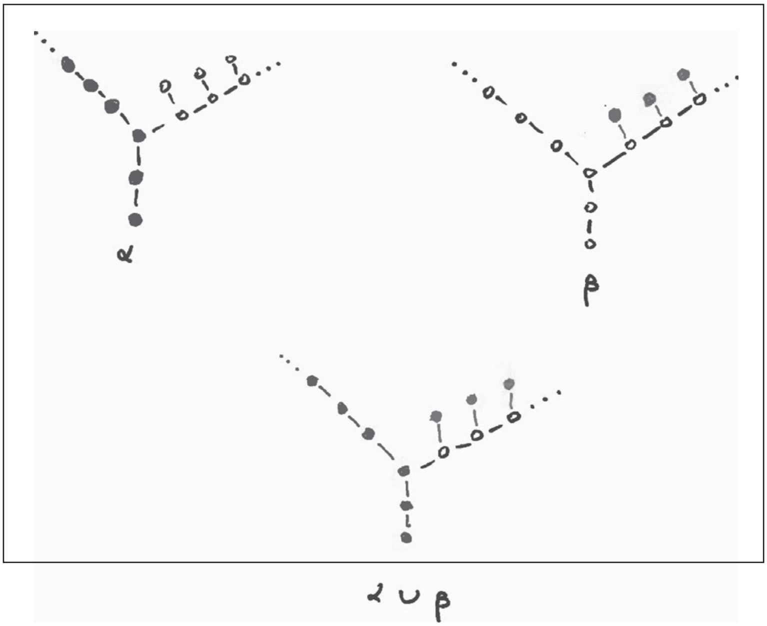

Fig. 5. Two sets of nodes α and β. Both sets are convex, prolific and past-complete, but their union is not convex.

like the stem events and some of the other events we have been using as illustrations. In attempting such an extension or 'prolongation' of μ, it is natural to think in terms of approximations, or more formally of limits. Let $A \subseteq \Omega$ be some event A outside the initial domain. In order to define $|A\rangle \equiv \mu(A)$ as a limit, we aim to identify a sequence $A_n \in \mathfrak{S}_n$ of 'best approximations to A' and hope that the corresponding measures $|A_n\rangle \in \mathfrak{H}$ also converge. If they do, we would then take their limit in $\mathfrak{H}$ to be the measure of A:

$$|A\rangle = \lim_{n\to\infty} |A_n\rangle \, .$$

Notice here that in attempting to define $|A\rangle = \mu(A)$ in this way, we have relied on two independent notions of convergence:

(i) the purely set-theoretic convergence of A_n to A in the sense of Section 5; and
(ii) the topological convergence of the measures $|A_n\rangle$ to $|A\rangle$ in Hilbert space (say in the norm topology or perhaps the weak topology).

We might ask whether the first notion is really needed, given that the extension theorems of ordinary measure theory do without it and just rely on the measure μ itself. Would it be possible to proceed similarly here? Unfortunately, this looks doubtful, even though the vector $|A_n\rangle$ carries a certain amount of information about the event A_n (but a very limited amount since, owing to quantal interference, very different events can share the same vector measure).

In the ordinary setting, where μ is real and positive, it defines a distance on the space of initially measurable sets modulo sets of measure zero such that two events $A, B \in \mathfrak{S}$ are

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

close when $\mu(A+B)$ is small. Extension of the measure then corresponds to completion of the metric space thereby defined (Kolmogorov and Fomin 1961). Quantally, however, an event of small or zero measure is not negligible in the same way as it is classically, because of interference. Thus, if we tried to use the norm of $|A+B\rangle$ as a distance, it would not even obey the triangle inequality. For example, for three disjoint events A, B and C, as in the 3-slit experiment of Sorkin (1997; 2007), we have $|A+B\rangle = |B+C\rangle = 0$ but $|A+C\rangle \neq 0$. Similarly, trying to quotient the event algebra by the events of measure zero would yield nonsense; it can even happen that all of Ω is covered by events of measure zero (Sorkin 2012). To establish an association between a vector $v = \lim |A_n\rangle$ and a definite event A in Ω, it appears that we need an independent notion of convergence like the one introduced in Section 5 and developed in Section 6.

Accepting this apparent necessity, let us investigate how a limiting procedure might go in the important case of an open event $E \in \bigvee \mathfrak{S}$. In doing this, we will employ the canonical approximants $E_n \in \mathfrak{S}_n$ 'of type C1' for E, these being the simplest to work with and probably the first to suggest themselves for most people:

$$E_n = \bigcup \{Z \in \mathfrak{Z}_n \,|\, Z \subseteq E\}. \tag{10}$$

(this is just (7) with E in place of A). As we know, there is no guarantee in general that the corresponding vectors $|E_n\rangle$ will converge, but when they do, we would like to regard E as 'measurable' and to associate with it the measure $|E\rangle = \lim_n |E_n\rangle$. We will illustrate this procedure using the two-site and three-site hoppers, but we first consider whether or not our criterion of convergence is adequate as it stands, or whether it needs to be strengthened.

Recall in this connection that we have already defined a second sequence of approximants for E 'of type C2', which are related to the first ones by past-completion of the corresponding node-sets α. Would these approximants have led to the same set of measurable open events and to the same measures for them? A second question concerns compatibility with the Boolean operations, some form of which is needed if the extended measure is to be additive on disjoint events. Consider, for example, the disjoint union $G = A + B$ of two open events A and B, and let G_n, A_n and B_n be the corresponding C1 approximants. If $|A_n\rangle + |B_n\rangle = |G_n\rangle$ held automatically, it would follow immediately that $|A\rangle + |B\rangle = |G\rangle$, as desired. But plainly this is not automatic because $G = A \cup B$ can include cylinder sets Z that are not included separately in either A or B (see figure 6). We would like the contribution from such Z to go away in the limit $n \to \infty$, and we would also like any mismatch between our C1 and C2 sequences to go away as well. These two requirements turn out to be closely related.

Let us examine the difference

$$|G_n\rangle - |A_n + B_n\rangle = |G_n\rangle - |A_n\rangle - |B_n\rangle \tag{11}$$

more closely. In light of equation (10), this difference is just

$$\sum \{|Z\rangle \,|\, (Z \subseteq A+B)(Z \nsubseteq A)(Z \nsubseteq B)\},$$

but the Z here are not arbitrary cylinder sets. Rather, we claim that any Z that contributes to the above sum is special in that its overlaps with A and B are clopen (which implies

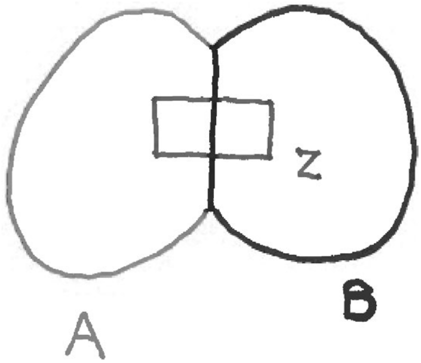

Fig. 6. A cylinder set Z contributing to the difference (11).

that, within Z, the discrepancy disappears entirely in a later approximation). By the next lemma, something similar holds for the difference between the C1 and C2 approximants to any individual open set A.

Definition 7.1. The cylinder set Z *straddles* the event A if it meets both A and its complement and if $Z \cap A$ is clopen. In symbols, $0 \neq ZA \neq Z$ and $ZA \in \mathfrak{S}$.

Remark 7.2. Z straddles A if and only if it straddles the complement $\Omega \backslash A$.

To see why the Z contributing to equation (11) are 'straddlers' in this sense, it suffices to observe first that $ZA \equiv Z \cap A$ is open because both A and Z are open, and second that ZB is therefore closed, having the form closed\open: $ZB = Z \backslash ZA$. By symmetry, both ZA and ZB are consequently both open and both closed.

Lemma 7.3. Let A be an open event and α be the corresponding node-set of type C1. If past-completion adds $Z \in \mathfrak{Z}$ to α, then Z straddles A, and conversely.

Proof. Since A is open, α is upward-closed, while for any α at all, it is true that $\beta = \{x \in \mathfrak{T} \,|\, x \prec \alpha\}$ is downward closed. Hence, the difference $\beta \backslash \alpha$, the exclusive past of α, is also downward closed; that is, it is a subtree of $\mathfrak{T}$. We now suppose, without loss of generality, that $Z \in \beta \backslash \alpha$. By definition, past-completion will adjoin Z to α if and only if no path originating at Z can remain within the subtree $\beta \backslash \alpha$. (In this case, it cannot leave $\beta \backslash \alpha$ and return to it later because subtrees are convex.) But this means that the portion of $\beta \backslash \alpha$ above Z is actually finite (by the infinity lemma). Consequently, for n sufficiently big, every descendant of Z is either in α or below no node of α. Translated into the language of subsets of Ω, this says that at a sufficient degree of refinement n, every cylinder set $Z' \in \mathfrak{Z}_n$ and within Z is either fully included within A (the former alternative) or disjoint from A (the latter alternative). Now suppose that past-completing α does adjoin Z to it. Then ZA is the union of the Z' belonging to the first family and is therefore clopen, straddling A. Conversely, if ZA is clopen, then it is a union of cylinder sets $Z' \in \mathfrak{Z}_n$ for some n. □

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

In view of these results, we can to some extent deal with the issues raised above by provisionally adding to our criterion of convergence the requirement that any 'straddling' cylinder sets contribute negligibly in measure as $n \to \infty$.

Definition 7.4. Let $A_n \in \mathfrak{S}_n$ be a sequence of clopen events. We then say that this sequence is *evenly convergent* with respect to μ if the following conditions hold:

(i) $A = \lim A_n$ for some $A \subseteq \Omega$.
(ii) $|A\rangle = \lim |A_n\rangle$ for some $|A\rangle \in \mathfrak{H}$.
(iii) $(\forall \varepsilon > 0)(\exists N)(\forall n > N)\ \sum \|\ |Z\rangle \| < \varepsilon$, where the sum ranges over all $Z \in \mathfrak{Z}_n$ that straddle A.

In view of the previous lemma, condition (iii) implies immediately that the C1 and C2 approximants for any open event E yield equivalent results. As desired, it also gives rise to additivity on disjoint open events, as we will see in the next theorem. In the statement of the theorem, the canonical approximants may for definiteness be taken to be those 'of type C1'. As we have just seen, exchanging any of them for type C2 would have no effect. (Notice also in the statement of the theorem that the Boolean sum $A + B$ coincides with the union $A \cup B$ when A and B are disjoint, that is, $AB = 0$.)

Theorem 7.5. Let A and B be disjoint open events and let A_n and B_n be their canonical approximating sequences, with G_n being the canonical approximating sequence for $G = A + B$. If the first two sequences are evenly convergent then the third is also, and the measures add, that is,

$$|A\rangle + |B\rangle = |G\rangle.$$

Proof. It will be convenient to work with the canonical sequences given by (10) since for them the approximants to A and B will be disjoint. In the above definition of an evenly convergent sequence, we need to establish conditions (i)–(iii) with A replaced by G:

(i) $G = \lim G_n$, is true by construction.
(ii) We must verify that $|G\rangle = \lim |G_n\rangle$ with $|G\rangle = |A\rangle + |B\rangle$. We have already learned in connection with equation (11) that $|G_n\rangle - |A_n\rangle - |B_n\rangle$ is the sum of the measures $|Z\rangle$ of all those cylinder sets $Z \in \mathfrak{Z}_n$ that straddle both A and B. But this sum can be made arbitrarily small by choosing n big enough since, by hypothesis, the sequence A_n is itself evenly convergent. In more detail:

$$\begin{aligned} \|\ |G_n\rangle - |A_n\rangle - |B_n\rangle \| &= \| \sum \{|Z\rangle \mid Z \text{ straddles both } A \text{ and } B\} \| \\ &\leqslant \sum \{\|\ |Z\rangle \| \mid Z \text{ straddles both } A \text{ and } B\} \\ &\leqslant \sum \{\|\ |Z\rangle \| \mid Z \text{ straddles } A\} \end{aligned}$$

and this $\to 0$ as $n \to \infty$.
Therefore

$$\begin{aligned} \lim |G_n\rangle &= \lim(|A_n\rangle + |B_n\rangle) \\ &= \lim |A_n\rangle + \lim |B_n\rangle \\ &= |A\rangle + |B\rangle, \end{aligned}$$

as required.

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

(iii) We need to verify that the G_n themselves fulfill the third condition for being evenly convergent. To that end, we will demonstrate that any cylinder event $Z \in \mathfrak{Z}_n$ that straddles G also straddles either A or B. The total norm of the straddlers of G will thus be bounded by the sum of the bounds for A and B, both of which go to zero as n goes to ∞, which will complete the proof.
Suppose, then, that Z straddles

$$G = A + B = A \sqcup B,$$

where the symbol '$\sqcup$' denotes the union of disjoint sets. We then have that

$$ZG = Z(A \sqcup B) = ZA \sqcup ZB.$$

By definition, Z meets $A + B$, so suppose it meets A, that is, $ZA \neq 0$. Now ZA is obviously open since both Z and A are open. It is also closed, being the difference of the clopen set $ZG = ZA \sqcup ZB$ and the open set ZB. Hence Z straddles A if it meets A at all, and in general it will straddle either A or B, as we set out to prove. □

This theorem takes a first step toward arranging additivity of the extended measure, but disjoint open events are, of course, a special case. More generally, we would like to have similar theorems covering, say, arbitrary events in $\mathfrak{R} \bigvee \mathfrak{S}$ (not just open events) and arbitrary Boolean operations (not just disjoint union). For example, it is easy to establish for any open event E that $|E\rangle$ is defined if and only if $|\Omega\backslash E\rangle$ is defined, and that then $|\Omega\backslash E\rangle + |E\rangle = |\Omega\rangle$. To what extent such results can be obtained in general remains to be investigated.

7.1. *Examples*

Our old friend the return event R can serve to illustrate some of the definitions we have made. Let us start with the 2-site hopper, in which case $R' = \Omega\backslash R$, the event of 'non-return', consists of the single history, $(0, 0, 0, \cdots)$. As we know, R itself is topologically open, and, correspondingly, $\Omega\backslash R$ is closed, as can be seen directly from the fact that it is the limit of a decreasing sequence of clopen events of the form $R'_n = \text{cyl}(0, 0, 0, \cdots, 0)$, these being our canonical approximants for R'. In this case, no cylinder event in $\mathfrak{Z}_n$ straddles R'_n since it is itself a cylinder event. So, to check that $|R'\rangle$ is well defined, we just have to check that the sequence $|R'_n\rangle$ converges. In fact, it converges trivially to 0, since $\|\, |R'_n\rangle \,\| = (1/2)^{n/2}$. Thus, $|\Omega\backslash R\rangle = 0$ and non-return is *precluded* for the two-site hopper (Gudder and Sorkin 2012). Taking complements then shows that $|R\rangle$ is defined and has the value $|R\rangle = |\Omega\rangle$.

In the context of the three-site hopper, the events of return and non-return become much more interesting. Classically, non-return 'almost surely' does not occur in a finite lattice, that is, its measure vanishes. Moreover, this conclusion follows independently of whatever initial conditions we care to assume. But what will we find quantally? More generally, what will we find if, instead of asking whether the particle visits site 0, we ask whether it visits site 1 or 2? By symmetry, these questions become equivalent if we generalise our initial condition to admit different starting sites. Let us therefore consider

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

(still more generally) an initial condition in which each possible initial location contributes its own complex amplitude $\psi_0(j)$, $j = 0, 1, 2 \in \mathbb{Z}_3$. The measure of a cylinder set of trajectories can then be derived from the 3-site analogue of equation (2), generalised to allow for an arbitrary initial position x_0 , and with an additional factor of the initial amplitude $\psi_0(x_0)$ thrown in:

$$v_y = (U^{-n})_{yx_n} U_{x_n x_{n-1}} \cdots U_{x_2 x_1} U_{x_1 x_0} \psi_0(x_0).$$

For the event of non-return, we must sum this expression over all trajectories x_j such that $x_j \neq 0$ for all $j > 0$. The resulting vector of components v_y is then evidently given by a matrix product of the form $(U^{-n})V^n\psi_0$, where the matrix V is just the matrix U with its first row set to zero. We can also set the first column of V to zero if we re-express v as $(U^{-n})V^{n-1}\psi_1$, where ψ_1 is just $U\psi_0$ with its first entry set to zero. In this way, V becomes effectively a 2×2 matrix.

Recall now that for three sites, we have (with $\omega = 1^{1/3}$)

$$U = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 & \omega & \omega \\ \omega & 1 & \omega \\ \omega & \omega & 1 \end{pmatrix}$$

and thus also

$$V = \frac{1}{\sqrt{3}} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & \omega \\ 0 & \omega & 1 \end{pmatrix}.$$

For these matrices, the powers of U and V can be evaluated in essentially the same way by writing U or V as a linear combination of orthogonal projectors. Taking U as an example, we obtain, by adding and subtracting a multiple of the identity matrix to U,

$$U = \lambda(1 - P) + \sigma P,$$

where

$$P = \frac{1}{3} \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}$$

and

$$1 - P = \frac{1}{3} \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix}$$

with $\lambda = (1 - \omega)/\sqrt{3}$ and $\sigma = (1 + 2\omega)/\sqrt{3}$.

In the same way, defining

$$Q = 1/2 \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$$

(or more correctly as the 3×3 matrix with this as its lower right-hand corner), we can obtain V in the form,

$$V = \lambda(1 - Q) + \rho Q,$$

where

$$Q = \frac{1}{2}\begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}$$

and

$$1 - Q = \frac{1}{2}\begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & -1 \\ 0 & -1 & 1 \end{pmatrix},$$

with $\rho = -\omega^2/\sqrt{3}$. It follows immediately that

$$\begin{aligned} U^n &= \lambda^n(1-P) + \sigma^n P \\ V^{n-1} &= \lambda^{n-1}(1-Q) + \rho^{n-1}Q. \end{aligned}$$

Noting that $|\rho| = 1/\sqrt{3} < 1$, while $|\lambda| = |\sigma| = 1$, we can see that in the limit $n \to \infty$, we can drop the second term in V without affecting $|R'\rangle$, that is, without affecting whether the sequence of approximations $|R'_n\rangle$ converges, or what it converges to. And noting further that $P(1-Q) = 0$, we see that we can also drop that term in the product $U^{-n}V^{n-1}$, leaving the simple asymptotic form,

$$\begin{aligned} U^{-n}V^{n-1} &\sim \lambda^{-n}(1-P)\lambda^{n-1}(1-Q) \\ &= (1/\lambda)(1-P)(1-Q). \end{aligned}$$

We thus obtain, modulo an exponentially small correction,

$$|R'_n\rangle = \frac{1}{\lambda}(1-P)(1-Q)\psi_1 \ .$$

This formula leads to a somewhat odd conclusion. With our original initial condition that the particle begins at 0, the components of ψ_1 are just the last two entries of the first column of U, namely $(\omega/\sqrt{3})(0,1,1)$, which evidently belongs to the kernel of $1-Q$. Hence the event of non-return is again precluded: $|R'\rangle = 0$; and once again $|R\rangle = |\Omega\rangle = (1,0,0)$. At first sight this result might appear to confirm our classical intuition, but, in fact, it appears to be a coincidence, at least if we take the 3-site hopper as typical, since $|R'\rangle$ does not vanish for almost any choice of initial amplitudes other than $(1,0,0)$! In particular, if the particle starts at site 2 instead of site 0, the event that it fails to visit site 0 has the non-zero vector measure

$$|R'\rangle = (1/3, -1/6, -1/6).$$

(The quantal measure of this same event in the sense of Sorkin (1994) is then $\langle R'|R'\rangle$, or $1/6$.) The vector measure of the complementary event that the particle does visit 0 is then

$$|R\rangle = |\Omega\rangle - |R'\rangle = (-1/3, 1/6, 7/6).$$

Our analysis of the 3-site case tacitly used the fact that the events R and R' are free of straddling cylinder sets, for the same reason that stem events are. Convenient though this is, it means that our example fails to illustrate condition (iii) in our definition of an evenly convergent sequence. It would be good to work out an example where (iii) does come into play, since doing so could indicate whether that condition is a reasonable one to have added, or whether, on the contrary, it tends to rule out events that we would want to include.

It would also be good to work out some physically interesting instances of our approximation procedure in the causal set case. We might begin, for example, with the event 'originary' for the relatively simple dynamics of complex percolation.

8. Epilogue: does physics need an actual infinity?

Does the description of nature require actual infinities? Or is a truly finitary physics possible, in which infinite sets would figure only as potentialities?

Inasmuch as the theories to which we have grown accustomed make heavy use of real numbers, they thereby presuppose an actual infinity of cardinality $\aleph_1$, as emphasised in Isham (2002). In itself, however, this seems more a matter of convenience than of principle, since we could imagine making do with rational numbers of a very fine but finite precision that could be made still finer as the need arose – in other words a potential infinity.†

The other prominent continuum in current physics is, of course, space–time. Non-relativistically, we could again imagine circumventing the actual infinities that continuous space and time seem to imply, but when it comes to relativistic field theories, the new requirement of *locality* appears to force strict continuity on us. Perhaps we could get by with only $\aleph_0$ points, say points with rational coordinates, but even that would still be an actual infinity.

Quantum gravity raises all these questions anew, of course. String theory and loop quantum gravity both presuppose background continua, at least in their current formulations. Causal dynamical triangulations and the 'asymptotic safety' approach retain locality and presuppose the same type of continuum as classical gravity, albeit not as background.

With causal sets, the situation seems more fluid. On the one hand, they transcend locality, but on the other hand they still maintain *covariance* in the sense of label-invariance, and that brings with it an 'infrared' infinity, as discussed earlier. An important new feature, however, is that now the infinity is in some sense pure gauge: we need it only because we have introduced both an auxiliary time parameter and a space of 'completed causets' in order to give a precise meaning to the concept of sequential growth. Could it be that a manifestly covariant formulation of growth dynamics could dispense with this 'last remaining infinity'? As we are limited to measure theoretic tools inherited from the classical theory of stochastic processes, we seem to lack the technical means to ask the question properly. As things stand, we can acknowledge at a minimum that being able to refer to completed causets is very convenient, even if it ultimately turns out not to be physically necessary. Note also that the cardinality of a completed causet, though not finite, is reduced to that of the integers. On the other hand, the associated sample-space Ω still has the cardinality of the continuum.

Based on this evidence, we could perhaps agree that physics is tending toward more finitary concepts, even if it has not genuinely reached them yet. In particular, even if

† In writing '$\aleph_1$', I have adopted the continuum hypothesis, $\aleph_1 = 2^{\aleph_0}$, for ... notational reasons.

causal sets are implicitly free of actual infinities, the available mathematical tools do not let us express this fact clearly. It may be that some of the tools we seem to lack will arise naturally in the course of attempts, like those above, to extract well-defined generalised measures from quantal path-integrals and path-sums.

Appendix A. Some symbols used in the paper (in approximate order of their appearance)

Ω (the sample-space or space of histories), $\Omega^{physical}$, Ω^{gauge}, $\Omega(n)$

$0 \subseteq \Omega$ (the empty subset)

$\mathrm{cyl}(c)$ (the cylinder event corresponding to the truncated history c)

$\mathfrak{Z}$ (the semiring of cylinder events), $\mathfrak{Z}_n$

$\mathfrak{S} = \mathfrak{R}\mathfrak{Z}$ (the Boolean algebra generated by $\mathfrak{Z}$ = the finite unions of cylinder sets)

$\mathfrak{S}_n$

$\mathfrak{T}$ (the tree of truncated histories), $\mathfrak{T}_n$

$\mathbf{1}^z \equiv \exp 2\pi i z$

$\bigvee \mathfrak{S}$, $\bigwedge \mathfrak{S}$, $\bigvee \bigwedge \mathfrak{S}$

$S(\alpha)$, $\widetilde{S}(\alpha)$

$\lim$, Lim, $\liminf$, $\limsup$

$\prec$

$\mathbb{P}$, $\overline{\mathbb{P}}$

$A \sqcup B$

$|Z\rangle = \mu(Z)$

Acknowledgments

I would like to thank Sumati Surya for numerous corrections and suggestions for improving the clarity of the manuscript.

References

Bombelli, L., Lee, J., Meyer, D. and Sorkin, R. D. (1987) Spacetime as a Causal Set. *Physical Review Letters* **59** 521–524.

Brightwell, G., Dowker, F., García, R. S., Henson, J. and Sorkin, R. D. (2003) 'Observables' in Causal Set Cosmology. *Physical Review D* **67** 084031. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

Brightwell, G., Dowker, H. F., García, R. S., Henson, J. and Sorkin, R. D. (2002) General Covariance and the 'Problem of Time' in a Discrete Cosmology. In: Bowden, K. G. (ed.) *Correlations, Proceedings of the ANPA 23 conference*, Alternative Natural Philosophy Association, London 1–17. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

https://doi.org/10.1017/S0960129511000545 Published online by Cambridge University Press

Dowker, F., Johnston, S. and Surya, S. (2010) On extending the Quantum Measure. (Available at arXiv:1007.2725 [gr-qc].)

Dowker, F. (2006) Causal sets and the deep structure of Spacetime. In: Ashtekar, A. (ed.) *100 Years of Relativity – Space–time Structure: Einstein and Beyond*, World Scientific 445–464.

Gudder, S. and Sorkin, R. D. (2012) Two-site quantum random walk (in preparation).

Henson, J. (2009) The causal set approach to quantum gravity. In: Oriti, D. (ed.) *Approaches to Quantum Gravity – Towards a new understanding of space and time*, Cambridge University Press 393–413. (An extended version of this review is available at arXiv:gr-qc/0601121v2.)

Isham, C. J. (2002) Some Reflections on the Status of Conventional Quantum Theory when Applied to Quantum Gravity. (Available at arXiv:quant-ph/0206090v1.)

Kechris, A. S. (1995) *Classical Descriptive Set Theory*, Springer-Verlag.

Kolmogorov, A. N. and Fomin, S. V. (1961) *Measure, Lebesgue Integrals, and Hilbert Space* (Translated by Natascha Artin Brunswick and Alan Jeffrey), Academic Press.

Martin, X., O'Connor, D. and Sorkin, R. D. (2005) The Random Walk in Generalized Quantum Theory. *Physical Review D* **71** 024029. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

Moschovakis, Y. N. (2009) *Descriptive Set Theory* (second edition), American Mathematical Society.

Pearle, P. (1973) Finite-Dimensional Path-Summation Formulation for Quantum Mechanics. *Physical Review D* **8** 2503–2510.

Rideout, D. P. and Sorkin, R. D. (2000) A Classical Sequential Growth Dynamics for Causal Sets. *Physical Review D* **61** 024002. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

Sorkin, R. D. (1994) Quantum Mechanics as Quantum Measure Theory. *Modern Physics Letters A* **9** (33) 3119–3127. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

Sorkin, R. D. (1997) Quantum Measure Theory and its Interpretation. In: Feng, D. H. and Hu, B.-L. *Quantum Classical Correspondence: Proceedings of the* 4th *Drexel Symposium on Quantum Nonintegrability*, International Press 229–251. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

Sorkin, R. D. (2005) Causal Sets: Discrete Gravity (Notes for the Valdivia Summer School). In: Gomberoff, A. and Marolf, D. (eds.) *Lectures on Quantum Gravity: Series of the Centro De Estudios Científicos), proceedings of the Valdivia Summer School*, Plenum 305–328. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

Sorkin, R. D. (2007) Quantum dynamics without the wave function. *Journal of Physics A: Mathematical and Theoretical* **40** 3207–3221. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

Sorkin, R. D. (2012) Logic is to the quantum as geometry is to gravity. In: Murugan, J., Ellis, G. F. R. and Weltman, A. (eds) *Foundations of Space and Time*, Cambridge University Press. (Preprint available at `http://www.perimeterinstitute.ca/personal/rsorkin/some.papers/`.)

Surya, S. (2011) Directions in Causal Set Quantum Gravity. (Available at arXiv:1103.6272v1 [gr-qc].)

Varadarajan, M. and Rideout, D. (2006) A general solution for classical sequential growth dynamics of Causal Sets. *Physical Review D* **73** 104021.