Recurrence Recharged

Bryan W. Roberts

Draft: April 21st, 2007
Comments welcome, but please do not quote.

Abstract

Extensions of eternal recurrence are developed, and used to probe the structure of indeterministic systems.

"This life, as you live it now and have lived it, you will have to live again and again, times without number."

F. Nietzsche. *The Gay Science* section 341.

0. Introduction

Eternal recurrence is perhaps often viewed as the fanciful suggestion of existentialists with little better to do. But the received view of eternal recurrence may be extended in ways that are both respectful of our intuitions and fruitful in other philosophical analyses.

Our strategy will be to consider examples of systems that are recurrent according to a particular theory of physics. The first section of this paper begins by developing three successively more general notions of recurrence, which are dubbed *simple recurrence*, *symmetric recurrence*, and *quantifiable recurrence*. The first is perhaps a familiar notion of recurrence, and the others humble generalizations. The second section turns an eye toward the philosophical analysis of indeterminism, and introduces two useful results in this study: Norton's (2006) Dome, and Belnap's (1992, 2003) theory of Branching Space-Time. The third section finally shows how eternal recurrence can be used to probe the structure of an indeterministic system. One result of this section is the construction of a recurrent system that is indeterministic. This 'branching' structure of this system is then explained and analyzed in the language of Branching Space-Time.

1. Extending the Received View of Recurrence.

Fuzzy Intuitions. Cutting Through the Fuzz. Extending the Received View.

1.1. Fuzzy Intuitions about Recurrence

The intuitive idea behind recurrence is that our world consists of a cyclic ordering of events. Put another way, a recurrent world has the property that every sequence of events will inevitably repeat. But there is much hidden behind this fuzzy intuition, which can be illuminated through the precise characterization of 1) an orderings of events, and 2) the notion of repetition. Indeed, the exploration of these two parameters ought to be central to any modern development of recurrence.

Yet the received view of recurrence, it appears, sets these two parameters only to their simplest values: it assumes a linear ordering of events, and characterizes repetition as translation symmetry on that ordering. The view sets out by adopting the so-called doctrine of Laplacian determinism, which states that any given state of the world A is inevitably followed by a determinate ordering of events leading to a later state B. If B turned out to be indistinguishable from A, then it would follow that these events would continue cycling eternally. Taking a deterministic ordering of events to be a linear one, the 'cycles' may then be informally characterized as forward and backward translation invariance on that ordering. (In fact, under certain conditions, the existence of such periodicities can be proved on the basis of Laplacian determinism, as Earman (1986,

§VIII) has shown.) Examples of what I call the 'received view of recurrence' may be found in van Fraassan (1970), Chapman (1982), Earman (1986)[1], and Earman (1995).

1.2. Cutting Through the Fuzz

Let me propose one precise way to characterize this view of recurrence by defining time-translation symmetry more rigorously. One way to do this is to imagine the real numbers as a copy of our linear ordering of events, which we might call O. To sketch the procedure first: we introduce translation symmetry onto the Reals through the definition of an equivalence relation '=' in which $a = b$ iff $a \sim b$ modulo n. This equivalence is interpreted as physical indistinguishability of events in a system. Time translation symmetry is then identified with the existence of such an equivalence relation, and simple recurrence is defined as the property of time-translation invariance on a linear ordering. Here is the precise procedure:

Think of each event in a linear ordering O as complete description of the spatial properties of some isolated system, whose behavior is completely described by a physical theory T. (For example, one might write down the Lagrangian function for this system and take each element of the pre-image to be an event.) If T has a finite number of degrees of freedom, then we can introduce an equivalence relation $=_T$ that indicates the 'physical indistinguishability' of events in O. In other words, $e_1 =_T e_2$ whenever e_1 and e_2 are indistinguishable according to the theory T[2]. Now introduce a time function t: O → **R**,

□

[1]Earman (1986) also suggests that one might study the recurrence of certain probability states on stochastic processes. This kind of recurrence is not discussed in the present work, as our goal is to argue for a more general kind of recurrence that can occur even in the absence of a probability space.

[2]There are of course physical theories with infinitely many degrees of freedom in which the existence of such an equivalence relation is not guaranteed; quantum field theories are the most notorious examples. However, it might still be possible to consistently introduce a relation representing physical

where t is an order-preserving bijection and **R** is the real number field. We may then propose the following definition of time-translation invariance.

Definition 1. A linear ordering of events O is *time-translation invariant*[3] whenever there exists an order preserving bijection t: O → **R** and some $x \in$ **R** such that for all e_1, e_2 in O, if $t(e_1) = t(e_2) \mod x$, then $e_1 =_T e_2$.

I have already suggested that a linear ordering of events admitting time-translation symmetry is the simplest of many possible kinds of recurrence. So in anticipation of more complex cases, let us call this property *simple recurrence*:

Definition 2. A system is *simply recurrent* if it may be characterized as a linear ordering of events that is time-translation invariant.

A toy model of a simply recurrent system may be constructed by taking a universe that consists only of a spherical mass of unit radius, which rotates constantly about some axis at one unit per second, for all of time.

Proposition 1: The spinning-sphere universe is a simply recurrent system.

Proof: For definiteness, adopt classical mechanics and a polar coordinate system. Take an event in the spinning sphere universe to be a triple (r, θ_0, ψ), for some fixed value θ_0, such that all of space is spanned by varying r and ψ. Then a linear ordering is

indistinguishability in these theories. And in any case, nearly all well-confirmed physical theories have only finitely many degrees of freedom, and so we can always simply restrict our attention to these cases.

[3] That is, time-translation invariant *according to a theory T*. The latter part is tacitly assumed and omitted here and henceforth.

induced on the set of events by the ordering of the θ-values. Furthermore, any event (r, θ_0, ψ) is physically indistinguishable from (r, $\theta_0 + 2\pi n$, ψ), for all $n \in \mathbf{Z}$. So $t(r, \theta_0, \psi) = t(r, \theta_0 + 2\pi n, \psi)$ is provides a time-translation invariance according to Definition 1, and therefore this system is simply recurrent by Definition 2.

There are many other ways that simple recurrence can be defined. However, one requirement of our approach to recurrence is that the definition remain silent on the question of whether or not a recurrent temporal ordering is topologically open or closed. Informally, a recurrent system that is topologically open admits to a chain of events that keeps recurring all the time, while one that is topologically closed admits only a single chain of events that 'loops' back upon itself[4]. Grünbaum (1963, 197-208) has warned that if there are no attributes that distinguish between two events in a recurrent system, the two events must be identical by Leibniz's principle. However, Grünbaum has taken 'identical' here to mean 'corresponding to one point on a spacetime manifold.' In contrast, we have set out to characterize recurrence *according to a given theory of physics*, and thus are justified only taking 'identical' to mean 'physically indistinguishable' according to that theory. Most physical theories simply lack the language to distinguish between topological openness and topological closure in the sense of Grünbaum's worry. To introduce this distinction into the discussion would thus seem to introduce a philosophical problem where there should be none. This is an important lesson, which we shall return to in section 3.3.

1.3. Extending the Received View

[4]We follow Grünbaum (1963) in adopting the language of openness and closure. Earman (1995), in contrast, calls the former 'recurrent time' and the latter 'cycle time.'

Given the route on which we have set out, there are two ways in which the notion of recurrence can be extended. We can either 1) extend what counts as repetition, or 2) extend what counts as an ordering of events. In fact, we shall do both. First, we provide two successively more general notions of repetition. We shall then discuss which of these notions can be applied to a non-linear ordering of events. As we discuss in the next section, dropping the linearity requirement allows one to develop an account of the recurrence of indeterministic systems.

In the definition of simple recurrence, repetition is interpreted as time-translation invariance. To understand how this requirement might be relaxed while retaining a coherent notion of recurrence, let us return to the rotating sphere model described above. Instead of a unit sphere rotating with constant velocity, let us now imagine that the rotation of the sphere stops and starts in a predictable way. Suppose, for example, that the sphere completes one revolution and then stops for 2π seconds; it then completes another revolution and stops for 4π seconds. Suppose that it continues doubling the time between revolutions in this way. Such a system is certainly not time-translation invariant according to Definition 1. However, it is invariant under the function that translates by twice as much after each step. This suggests that *general* functional invariance might replace time-translation invariance in the definition recurrence. In both situations, definitions, the essential feature is that some phenomenon is guaranteed to recur again and again. So let us define our first more general notion of recurrence:

Definition 3. A linear ordering of events O is *1-symmetrically recurrent* if for all e in O, there is a function f: $O \times \mathbf{Z} \rightarrow O$ such that for all $i \in \mathbf{Z}$, $f(e, i) =_T e$.

We write the prefix to *1-symmetry* with the recognition that many more general kinds of symmetry can still be characterized. For example, we might have taken the domain of f to be $O \times \mathbf{R}$. Or, we could have given an algebraic characterization of symmetry, as (for example) the automorphisms of some group structure. However, as we show shortly, none of these characterizations seem to allow recurrence in systems that are indeterministic. So since our current goal is to characterize recurrence in indeterministic systems, we shall not attempt these developments here.

Instead, we achieve an altogether more general notion of recurrence by taking the last example of a rotating sphere to the extreme. Imagine that the number of seconds that the sphere remains at rest is changing erratically. For example, suppose each successive stop is for a finite but completely random period of time. This new sphere revolves once, then stops for four seconds; it revolves again, then stops for one second; it revolves again, and stops for thirty-seven seconds; and so on without pattern. The system is certainly not simply-recurrent according to Definition 2. Nor is it 1-symmetrically recurrent according to Definition 3. Nevertheless, this system shares one essential feature with both of these definitions: the sphere is guaranteed to revolve again and again, times without number. In other words, the rotation of the sphere is an event which, for any given time t, is guaranteed to occur somewhere in the future of t. This property characterizes what we shall call *quantifiable recurrence*:

Definition 4. A partial ordering of events is *quantifiably recurrent* if there exists some event e^* such that, for every e_1, there is an event $e_2 > e_1$ such that $e^* =_T e_1$.

Notice that in Definition 4, the requirement that our ordering of events be linear omitted. Both simple recurrence and 1-symmetric recurrence could not have been sensibly defined on anything but a linear ordering. The unfortunate side-effect is to restrict our attention to systems in which events evolve deterministically. Yet, as is well known, determinism fails not only in quantum mechanics, but in classical and general relativity as well. Examples are given by Earman (1986), Perez Laraudogoitia (1996, 1997), and Earman (1995). Thus, since our approach is to study recurrence according to physical theory, it is desirable to drop the requirement of a linear ordering of events.

Some reply that our goal ought to be to *purge* indeterminism from physics wherever possible, and restore more mild-mannered, linear orderings of events. However, this strategy has not yet been shown to be successful, and recent surveys (c.f. Earman (1986, 2004) suggest that there is little hope for the project. We shall therefore proceed with an eye toward indeterminism, and the hope that our account makes sense no matter how the debate is resolved.

In the next section, we examine a simple example of an indeterministic system to help guide our intuitions. We then adopt a precise formalism for discussing certain features of such systems. This will put us in a position to explore how the notion of quantifiable recurrence can be seen to apply to indeterministic systems such as this.

2. Classical Indeterminism.

Indeterminism Exposed. Concrete Example. Foundations in BST.

2.1 Indeterminism Exposed

Let us begin with an informal definition of what is meant by indeterminism according to a physical theory. We shall then see an example of a system that satisfies this definition, and proceed to precisely characterize the indeterminism that it exhibits.

> **Definition 5**. A system is *indeterministic* according to some theory of physics whenever a complete specification of initial and boundary conditions for that system, together with the laws of the theory, fails to uniquely determine the future time-evolution of that system.

This is one popular use of 'indeterminism,' but there are many others. Some readers may prefer to use a different term for Definition 5 and reserve 'indeterminism' for other uses; I trust such readers can substitute their favorite alternative (such as 'pathological initial value problem,' etc.) wherever necessary.

Of the many examples of classical indeterminism in the literature, we need only examine one of the simplest: Norton's (2006) dome. Norton's dome is an indeterministic system according to the theory of classical mechanics. Interestingly, a slight modification of the dome can result in a system that is both indeterministic and quantifiably recurrent. But in order to discuss this example of indeterministic recurrence, let us first review the pure indeterminism of Norton's dome.

2.2. Concrete Example: Norton's Dome

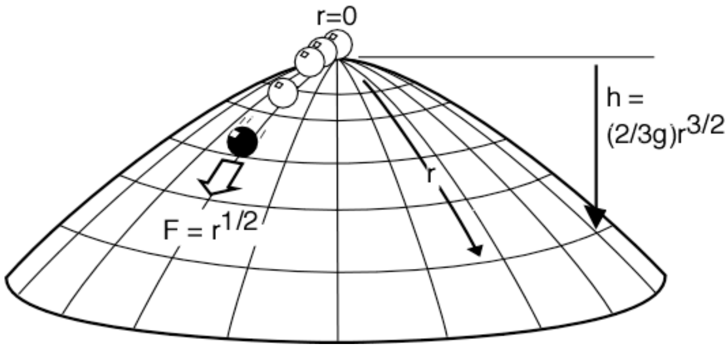

Figure 1: Norton's Dome. Illustration found in Norton (2006).

Norton (2006) presents a system in which a unit mass sits at the apex of a frictionless dome. The relation between the height *h* and the radius *r* is given by $h = \frac{2}{3g} r^{3/2}$ (see figure 1). One can easily show that the equation of motion describing the position of the unit mass on this surface is $\frac{d^2 r}{dt^2} = r^{1/2}$. Surprisingly, this equation admits an infinite family of solutions[5], indexed by the real numbers **R**:

$$
\begin{aligned}
r(t) &= \tfrac{1}{144}(t - T)^4 \qquad &\text{for } t \geq T \\
r(t) &= 0 \qquad &\text{for } t \leq T.
\end{aligned} \tag{1}
$$

Each of these solutions is interpreted as describing a ball at rest until some time $T \in \mathbf{R}$ when the ball spontaneously begins rolling down the dome. (Proofs of these facts may be found in Norton (2006).)

There is, of course, the additional (trivial) solution that r(t) = 0. However, the usual formulation of classical mechanics is not strong enough to choose this solution over

[5] As Malament (2006) has shown, the non-uniqueness of the solution to the initial value problem is ultimately due to the fact that the curvature of Norton's dome diverges at the apex.

any of the 'non-trivial' ones in which the ball rolls off the dome at an arbitrary time. Therefore, Norton's dome is an indeterministic system according to Definition 5.

Let us now examine Norton's dome in the context of a more precise formalism, that will allow us to sharpen our understanding of recurrence in the next example. A precise formalism for dealing with indeterminism has already been developed by Belnap (1992) and others; in this paper, we thus adopt their theory of Branching Spacetime (BST)[6].

2.3. Foundations in Branching Spacetime (BST)

BST is a theory with signature {S, <}, where S is a set of events and < is a partial ordering. Most of the central ideas make use of the following two definitions:

Definition 6. A *history* is a maximal, directed subset of S.

A history, as expected, is taken to represent the past and possible future of an event. In any interesting model of BST, there is more than one history:

Definition 7. A *choice point* for two histories is an element e in S that is common to both histories, such that every e* > e fails to be common to both.

The existence of choice points is the essential characteristic of an indeterministic system. That is, a choice point is a place where the future splits into two or more possible developments (such as 'measured spin up' and 'measured spin down').

[6] Strictly speaking, 'branching time' is sufficient for the simple examples we present here. However, we expose these ideas in the context of the more general theory of BST, in order to ease extension to more complex examples (e.g., systems admitting both indeterminism and non-locality).

The theory of BST is known to be readily and fruitfully satisfied by various indeterministic scenarios; for example, see Placek (2002, 2004), Belnap (2005), and Müller (2005). What is less well known is that BST also provides an illuminating framework in which to study recurrence in indeterministic scenarios. This latter application is the subject of the next section; in the remainder of the present section, we prepare for this study by showing how to describe Norton's Dome in the language of BST.

The dome is a Newtonian system; that is to say, it satisfies the axioms of Newtonian mechanics, and labels points by way of the usual Euclidean manifold $\mathbf{R}^3 \times \mathbf{R}$ with an ordering relation < defined on the 'time' dimension $\mathbf{R}$. Let us define a 'dome event' e_d to be a coordinate $(\boldsymbol{x}, t)$ in $(\mathbf{R}^3 \times \mathbf{R})$ specifying the position of the ball on the dome at time t[7]. The branching structure of the dome may then be specified using the following definition:

> **Definition 8**. A *dome history* h_d a maximally extended, directed set of events $\{(\boldsymbol{x}, t)\}$ on the dome that satisfies 1) the axioms of Newtonian mechanics, and 2) the initial conditions specifying the structure of the dome, such as the shape of the surface and the mass of ball.

[7] A more standard way of expressing the lifetime of the ball might be to write down a Lagrangian describing the system, together with the Euler-Lagrange equations. Then the pair (q, p), where $p = dx/dt$, would be sufficient to characterize any deterministic evolution. (For a more detailed treatment of this method, see Butterfield (2004, §3.3).) However, in an indeterministic system such as Norton's dome, a *family* of Lagrangians would have to be specified, one for each possible time-evolution. The precise specification of such a family is too complicated for the purposes of exposition, especially in the more complex Dome-In-Bowl system. Thus, let us represent events by position-time pairs instead.

As we saw in the last section, these specifications are satisfied by an infinite family. There is one history in which the ball remains at the apex for all of time, corresponding to the solution $r(t) = 0$ to the dome equations of motion. The rest of the histories each correspond to a solution in which the ball falls off the dome at time $T \in \mathbf{R}$, given by equation (1) above.

Thus, every event at which the ball is at the apex of the dome is a choice point, and two histories branch from each of these points: one in which the ball begins its descent, and another in which it does not. Furthermore, if we consider any history h_d in which the ball begins to fall at time T, every event at time $t \leq T$ is a choice point, while every event at time $t > T$ is not.

The branching structure, with the spatial dimensions suppressed, can be imagined as an infinitely long, infinitely dense comb. The spine of the comb is the (trivial) history corresponding to the solution $r(t) = 0$, in which the ball never falls off the dome. A 'tooth' branches off from each point on the spine of the comb, each tooth representing a different history in which the ball has fallen off the dome.

Notice that there is a single history in this system that is *recurrent*, namely, the history in which the ball remains at the apex of the dome for all of time. This history is simply recurrent according to Definition 2, since it is non-dynamical and thus trivially time-translation invariant. It thus meets our more general definitions of recurrence as well. However, a more interesting example would be one in which *every* history were recurrent in some sense, including the dynamical ones. In order to achieve an example of this kind, we need only make a slight modification to Norton's dome. The result will be an example of recurrence in an indeterministic system.

3. Indeterministic Recurrence.

Indeterministic Recurrence Exposed. Concrete Example. Constraints on Branching Structure.

3.1 Indeterministic Recurrence Exposed.

Our first two definitions of recurrence, *simple* and *1-symmetric* recurrence, were both properties of linear orderings. But intuitively, an initial value problem on such an ordering can never be non-unique, and therefore never be indeterministic by Definition 5. As a result, these types of recurrence are not ideal candidates in to study in indeterministic scenarios. The following proposition illustrates this more precisely.

Proposition 2. A linear ordering of events is deterministic.

Proof. We proceed in the language of BST. The only maximal directed subset of a linear ordering of events *O* is *O* itself. As a result, such orderings admit only one history (by Definition 6), and no choice points (by Definition 7). But an ordering of events that has no choice points is deterministic; therefore, *O* is deterministic.

On the other hand, quantifiable recurrence remains a promising possible property of indeterministic systems because it requires only a partial ordering of events. What we would like is a system that is dynamic, indeterministic, and quantifiably recurrent. One example of such a system is the Dome-In-Bowl.

3.2. Concrete Example: The Dome in a Bowl.

A quantifiably recurrent system that is indeterministic according to classical mechanics may be constructed by placing Norton's dome inside a large bowl and smoothing out the edges. (The resulting surface reminds one of a lemon-juicer.) The curvature of the dome portion of the surface is described by $h = \frac{2}{3g} r^{3/2}$, just like Norton's dome. But at some height h^* (the bottom of the dome), the surface begins to curve back up in the shape of a bowl. The Dome-In-Bowl is further required to have a finite radius, and a lip of the bowl portion that is higher than the apex of the dome.

When a ball is placed at the apex of the Dome-In-Bowl, the result is a system that is indeterministic in the same sense as Norton's dome: every event in which the ball is at the apex of the dome is a choice point. But unlike Norton's dome, the branching structure of the Dome-In-Bowl is much more complicated. The finite radius and sufficiently high lip of the bowl guarantee that in any history in which the ball falls off the dome, the ball will eventually come to a stop at some point at equal height on the bowl, and then reverse its path. A ball in such a reverse trajectory will arrive at rest at the apex of the dome in a finite amount of time (this is proved by the fact that classical mechanics is time-reversal invariant). It thus follows that in every history of the Dome-In-Bowl, we are guaranteed the existence of an infinitely repeating event, namely, the event in which the ball is at the apex of the dome. We state this with more precision in the following proposition:

Proposition 3. The Dome-In-Bowl is a quantifiably recurrent system.

Proof. Let τ be the amount of time it takes for the ball to roll away from the apex of the dome and back again. Let $e_2 > e_1$ be any two events on the Dome-In-Bowl. If e_2 is an event in which the ball is at the apex, then define $e^* = e_2$. If not, then there is an event $e^* > e_2$ in which the ball is at the apex (where the time between e_2 and e^* is less than τ). Therefore, Definition 4 is satisfied, and the Dome-In-Bowl is a quantifiably recurrent system.

As in the case of Norton's dome, there is one history in which the ball remains at the apex of the Dome-In-Bowl for all of time. This history is quantifiably recurrent in an uninteresting way, because it is non-dynamical. However, unlike Norton's dome, this is not the only recurrent history. Every history of the Dome-In-Bowl is quantifiably recurrent, and every history except the trivial one is dynamical. We thus have the unusual result that the ball in the Dome-In-Bowl is not guaranteed to be at a particular location at any given time t, but it is guaranteed to be at the apex at some time $t^* > t$.

3.3. Branching Structure

Let us finally observe two features about the branching structure of the Dome-In-Bowl as an indeterministic system, which make this a particularly unusual and interesting example. First, we show that the Dome-In-Bowl admits an infinitude of non-intersecting histories. Second, we show how Grünbaum's distinction between topological openness and topological closure leads to serious difficulties when applied to the Dome-In-Bowl, no matter which view we choose to adopt. I take this latter result to be further support toward my suggestion that Grünbaum's distinction between topological openness and topological closure can only lead to philosophical problems where there should be none. We make this first argument by giving an informal proof of the following proposition.

Proposition 4. The Dome-In-Bowl admits an infinitude of non-intersecting histories.

Proof. Let t_0 be some time at which the ball begins its descent down the dome, and let τ be the amount of time it takes to return again. Define h to be the history in which the ball falls off the apex at time $t_0 + n\tau$, for all $n \in \mathbf{Z}$. Intuitively, h is a history in which the ball falls off the dome 'as often as possible.' Now consider the history h_ε that is identical to h_1, except that it is translated forward by some small amount. More precisely, choose any arbitrary $\varepsilon \in (0, \tau)$, and define h_ε to be the history in which the ball falls off the apex at time $t_0 + \varepsilon + n\tau$, for every $n \in \mathbf{Z}$. Then since every choice point in the Dome-In-Bowl is an event at which the ball is at the apex, every choice point in h occurs at a time $t + n\tau$, for some $n \in \mathbf{Z}$. But none of these choice points are contained in h_ε, since choice points in h_ε occur ε later in time. Therefore, h and h_ε do not intersect. Furthermore, since ε was arbitrarily chosen, there are in fact infinitely many histories h_ε that do not intersect h, one for every point in the interval $(0, \tau)$.

One consequence of this is that questions about 'possibility' in the Dome-In-Bowl must be approached with some care. For example, there is a sense in which the question, 'is it possible that the ball fall off the dome at time $t + n\tau$, for all $n \in \mathbf{Z}$?' should be answered in the affirmative, since the desired history can clearly be constructed. However, as we have now seen, this response must be coupled with the caveat, 'but not if, for some ε, it is possible that the ball fall off the dome at time $t + \varepsilon + n\tau$, for all $n \in \mathbf{Z}$.' This modal muddle of an answer can be cleared up by switching to the language of BST: both h and h_ε are histories of the Dome-In-Bowl. However, there is no consistent set of histories of which both h and h_ε are members.

The question of whether or not the Dome-In-Bowl should be considered 'topologically open' or 'topologically closed' must be approached with even more care. Let us first suppose that the Dome-In-Bowl is topologically open. Then it follows there are 'distinct' events that are 1) physically indistinguishable, and 2) occur at the same time. To see this, let h_0 be the history in which the ball remains at the apex of the Dome-

In-Bowl for all of time, and let h be a history in which the ball falls at time t. Then in both h_0 and h, the ball be at the apex of the dome at time $t + \tau$, where τ is the amount of time it takes for the ball to complete its circuit. The theory of classical mechanics (or any other physical theory, for that matter) cannot distinguish h_0 and h at time $t + \tau$. However, if the branching structure of the Dome-In-Bowl is topologically open, then histories can never 'connect up again,' and so h_0 and h nevertheless contain distinct events at time $t + \tau$. This is an uncomfortable state of affairs, especially if one's conscience is burdened by Occam's razor, Leibniz's principle, or an empiricist streak. So let us take Grünbaum's advice, and consider the alternative.

Suppose that if two events in a recurrent system are indistinguishable, then they are in fact the same event. This will entail, for example, that the histories h_ε described above are topologically closed. However, it also entails the disagreeable property of 'backwards branching,' in the following sense.

Proposition 5. If the Dome-In-Bowl is topologically closed, then there is an event that admits incompatible events in its proper past.

Proof. Let e be an event in which the ball is *in motion* time t, in some history h. Let e_0 be the event in which the ball is *at rest at the apex* at time t, in the history h_0, where h_0 is the history in which the ball remains at the apex for all of time. Since a ball in motion in the Dome-In-Bowl will always return to the apex, there must be some $\varepsilon < \tau$ such that the ball is at rest at the apex in h at time $t + \varepsilon$. But the ball is also at rest at the apex in h_0 at time $t + \varepsilon$. These two events are physically indistinguishable and occur at the same time, so by topological closure, they are the same event. Call this event e^*. Then $e < e^*$ and $e_0 < e^*$, even though e and e_0 are clearly incompatible.

This 'backward branching' of histories, if possible, would be highly unusual. In particular, it would result in a non-unique past, and thus render the question, 'Did *e* occur in the past of *e**?' indeterminate. This state of affairs so rudely affronts one's intuitions that it is strictly prohibited in the theory of BST[8]. It thus appears equally undesirable to consider the branching structure of the Dome-In-Bowl to be topologically closed.

Having been pierced by two horns of a dilemma, let me make the following suggestion: perhaps the distinction between topologically open and topologically closed recurrent systems is not a distinction after all. A red flag should have been raised signaling that there is no empirical phenomenon that might distinguish between these two branching structures, nor is there a well-confirmed physical theory that can tell the difference. Another should have gone up when both alternatives led us into a muddle. I thus recommend that we retain our coherent account of indeterministic recurrence, and discard a distinction that was ill-motivated from the outset.

4. Conclusion

We have now provided three successively more general notions of recurrence, and have now discussed examples of each:

- A system is *simply recurrent* if it may be characterized as a linear ordering of events that is time-translation invariant.
- A system is *1-symmetrically recurrent* if it may be characterized as a linear ordering O in which, for every e1 in O, there is a function f:OxZ→O (where Z is

[8]In particular, 'no backwards branching' is an almost immediate consequence of the axiom of Prior-Choice. See Belnap (1992, 2003).

the set of integers) such that for all i in Z, f(e1, i) is physically indistinguishable from e1.

- A system is *quantifiably recurrent* if it may be characterized as an arbitrary ordering of events in which there exists some event e* such that, for every event e1, there is an event e2 > e1 such that e* is physically indistinguishable from e1.

The first definition represents the received view of recurrence, which was found to be altogether too narrow. The second definition extends the received view to include more general temporal symmetries beyond mere time-translation. The third led us to an analysis of indeterministic recurrence.

I hope to have shown that the notion of recurrence can be extended in fruitful ways beyond its traditional appraisal. In particular, the analysis of recurrence in indeterministic systems is one viable way in which the structure of indeterminism might be probed. However, further work toward understanding this structure remains to be done. For example, the 'consistent sets of histories' in the Dome-In-Bowl still remain to be carefully characterized. Such an analysis would seem to provide further insight into the branching structure of recurrent systems; however, these suggestions will have to await future development.

Bryan W. Roberts
Department of History & Philosophy of Science
University of Pittsburgh
bwr6@pitt.edu

References

Belnap, Nuel. 1992. Branching space-time. *Synthese*, 92:385-434.

Belnap, Nuel. 2003. Branching space-time, postprint January, 2003. Unpublished manuscript (17 Feb 2003), http://philsci-archive.pitt.edu/archive/00001003/

Belnap, Nuel. 2005. A Theory of Causation: *Causae causantes* (Originating causes) as Inus Conditions in Branching Spacetimes. *British Journal for the Philosophy of Science*, 56:221-253.

Butterfield, Jeremy. 2004. Between Laws and Models: Some Philosophical Morals of Lagrangian Mechanics. Unpublished manuscript (3 Sep 2004), http://arxiv.org/abs/physics/0409030

Chapman, T. 1982. *Time: A Philosophical Analysis.* Dordrecht, Holland: Reidel Publishing Company.

Earman, John. 1986. *A primer on determinism*. Dordrecth, Holland: Reidel Publishing Company.

Earman, John. 1995. *Bangs, crunches, whimpers and shrieks: singularities and acausalities in relativistic spacetimes.* New York: Oxford University Press.

Earman, John. 2004. Determinism: What We Have Learned and What We Still Don't Know. In *Freedom and Determinism*, pp. 21-46, Joseph Keim Campbell, Michael O'Rourke and David Shier (Eds.), Cambridge: MIT Press.

Grünbaum, Adolf. 1963. *Philosophical Problems of Space and Time*. New York: Random House.

Malament, David. 2006. Norton's Slippery Slope. Forthcoming in, *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, Vol. 2006.

Müller, Thomas. 2005. Probability Theory and Causation: a Branching Spacetimes Analysis. *British Journal for the Philosophy of Science* 56:487-520.

Nietzsche, Friedrich W. 2001. *The gay science: with a prelude in German rhymes and an appendix of songs*. Edited by Bernard Williams, translated by Joseline Nauckhoff. Cambridge: Cambridge University Press.

Norton, John. 2006. The Dome: An Unexpectedly Simple Failure of Detereminism. Forthcoming in, *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, Vol. 2006.

Perez Laraudogoitia, Jon. 1996. A beautiful supertask. *Mind*, New Series, 105:81-83.

———. 1997. Classical Particle Dynamics, Indeterminism and a Supertask. *British Journal for the Philosophy of Science*, 48:49-54.

Placek, Thomas. 2002. Partial Indeterminism is Enough: a branching analysis of Bell-type inequalities. In *Non-locality and Modality*, Jeremy Butterfield and Thomas Placek (Eds.), Dordrecht: Kluwer Academic (Nato Science Series, vol. 64).

Placek, Thomas. 2004. Screening-off Conditions in Bell's Theorem: a Branching Space-times Analysis. In *Logica Yearbook 2004*, pp. 243-354, Bihounek and Bilkova (Eds). Prague: Filosofia.

Tipler, Frank J. 1974. General Relativity and Eternal Return. In *Essays in General Relativity*, F. J. Tipler (ed.), New York: Academic Press.

van Fraassan, Bas C. 1970. *An Introduction to the Philosophy of Time and Space*. Hanover, MA: Random House.