Foundations of Discrete Causal Theory.

Ben Dribus

May 10, 2011

Introduction.

Discrete causal theory is an attempt to unify theoretical physics by replacing **spacetime** with a **discrete causal structure**. In discrete causal theory, the concept of causality is fundamental, while the concept of spacetime is derivative; spacetime is hypothesized to be the macroscopic manifestation of causal relations among elements of a discrete set. This point of view reverses the paradigm of **general relativity**, in which a given event *may* influence any event in a particular region of spacetime (its **future light cone**) and *may* be influenced by any event in another region of spacetime (its **past light cone**). This is illustrated in figure 1 below, which shows a region of flat two-dimensional spacetime on the left and a discrete causal structure called a **causal graph** on the right. Thus, spacetime geometry determines the scope of causality in relativity, while causality determines (emergent) spacetime geometry in discrete causal theory. The proposal that causality, together with the metric properties of spacetime, emerge from a single binary relation on an underlying set, is called the **causal metric hypothesis**. Discreteness and the causal metric hypothesis are the fundamental axioms of discrete causal theory.

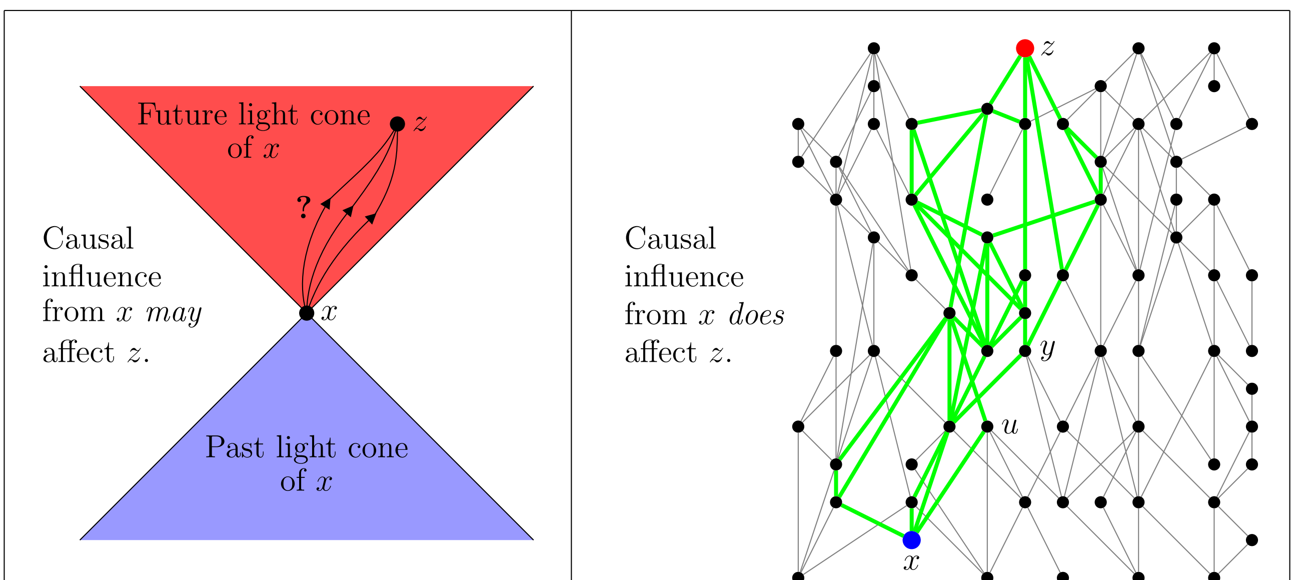

FIGURE 1

Discrete causal theory has the advantages of conceptual simplicity, **background independence**, and of course **discreteness**. It has the disadvantage of being in a primitive state of development, too primitive to be decisively predictive or falsifiable. Discrete causal theory has attracted less interest from physicists and mathematicians than other advanced theories such as **string theory** and **loop quantum gravity**.

Three important examples of discrete causal theory, in order of decreasing generality, are **causal graph theory**, **causal set theory**, and **causal dynamical triangulation**. These notes will focus on causal graph theory, but not exclusively. The research project developed here has three general goals: to lay down a solid foundation for discrete causal theory, to address some perceived shortcomings in existing versions of the theory, and to refine the theory to the point where it can be better compared to other theories and experimentally tested. The ultimate hope, of course, is that discrete causal theory can provide a **unified theory** of the fundamental forces of nature. Among the less grandiose but still highly ambitious specific goals of the project are a successful description of **quantum gravity** and a more satisfactory interpretation of the **foundations of quantum theory** than is currently available.

Overview of General Discrete Causal Theory.

The basic ingredients of discrete causal theory are **discrete sets**, **causality**, and **probability**. Given **causal structures** built from these ingredients, it remains to provide a means of describing and predicting their behavior. This is divided into two steps. First, one describes what types of behavior are possible; this is called **discrete causal kinematics**. Finally, one specifies a **dynamical law** for predicting what behavior actually occurs in specific situations, and uses this law to solve initial and boundary value problems; this is called **discrete causal dynamics**. The full development of these topics is a major undertaking and currently remains a work in progress, but at present I can at least offer a fairly clear outline of what is involved.

Discreteness.

The concept of a discrete set is familiar enough, but in fact several distinct **topological** and/or **metric** concepts come into play in discrete causal theory, and it is important to place these concepts in the proper context to one another. In the simplest possible terms, discreteness means that the theory has finite building blocks, called **elements**, which cannot be reduced to simpler objects. Ideally, as in causal graph theory and causal set theory, these elements have no **internal structure** at all; the structure of the theory comes entirely from specifying a **binary relation** called the **causal relation** on the underlying discrete set. The set together with the binary relation is called the **causal structure**. In practice, it is often easier to make contact with known physics by allowing some simplifying assumptions, and some versions of discrete causal theory allow for a certain amount of internal structure in the elements. For example, in the theory of causal dynamical triangulation, the elements are a special type of four-dimensional triangle called a **lorentzian simplex**. It is natural to question the meaning of the "finiteness" of the elements. Given a causal structure, metric properties are inferred from the properties of the causal relation in such a way that each individual relation is taken to be finite. In this sense there is a "finite separation" between any pair of elements of the discrete set, and this separation can be interpreted as "size" or "volume" of the elements. A common practice in causal set theory is to explicitly assign a "unit of volume" to each element of the underlying set, but there is no need to make this identification at the outset. A third topological idea that appears in discrete causal theory is the concept of **scale-dependent** or **atomic topologies**.

Ordinarily, the underlying set in a discrete causal theory is taken to be not only discrete, but **countable**. This allows the set to be embedded as a **discrete subset** (in the **relative topology**) of a **Euclidean space** (or sometimes a **Minkowski space**) for the purposes of visualization and illustration. The relations among the elements of the discrete set may then be represented by directed curves or line segments (Minkowski space provides its own "direction"), and the result is a concrete Euclidean or Minkowski **embedding** of a **directed graph**. More generally, any discrete causal structure, countable or uncountable, with no auxiliary structure, is an **abstract directed graph**. This is why causal graph theory is the most general form of pure discrete causal theory. For countable causal structures, Euclidean or Minkowski embeddings are very useful conceptually, but it is important to remember that they produce a misleading picture of the origin of metric structure. Indeed, in such an embedding, the **manifold structure** of the embedding space imposes a spurious metric structure on the causal structure. The paradigm of discrete causal theory is precisely the opposite: that the metric structure of physical spacetime is an emergent property of the binary relation on an underlying discrete set. This is discussed in more detail below.

There are several good reasons for preferring discrete theories over nondiscrete theories. First, discreteness is one of the fundamental characteristics of **quantum theory**. To be sure, not every aspect of ordinary **quantum mechanics** is discrete, but repeatedly over the last century discreteness has proven to be the key conceptual ingredient in solutions to both experimental and theoretical problems in microscopic physics. In **quantum field theory**, one begins with a continuous **classical field** (e.g. a **Maxwell** field or **Yang-Mills** field) and then carries out a process of **quantization** in which **elementary particles** (discrete entities) arise as the **quanta** of the classical field. General relativity describes spacetime itself as a continuous classical field, and although quantization of general relativity presents some unique challenges[1] compared to Maxwell theory and Yang-Mills theory, there is nothing particularly revolutionary about expecting discrete spacetime quanta to emerge from a successful theory of quantum gravity. In the spirit of some of the former successes of discreteness hypotheses, such as the

[1]The techniques used to quantize Maxwell theory or Yang-Mills theory lead to **nonrenormalizability** for general relativity. Some practitioners of loop quantum gravity believe that they have succeeded in quantizing general relativity by means of a different procedure called the **loop representation**, which was originally conceived as an approach to Yang-Mills theory but proved less successful for that purpose. Loop quantum gravity should certainly be taken seriously, but it suffers from some limitations. I will discuss loop quantum gravity briefly in the second part of the introduction, and the theory will make additional appearances for comparative purposes at various points throughout these notes.

analysis of the **black-body problem** that lead to the definition of **Planck's constant**, discrete theory holds promise for resolving some of the outstanding problems in microscopic physics involving the appearance of intractable infinities in calculations. The very presence of such infinities should be viewed as evidence of irrelevant information in the theory, and dubious nondiscrete structure is an obvious candidate for the source of the irrelevance.

Discreteness is sometimes viewed as "the alternative" to **continuum theories**, or theories involving manifolds. But of course discrete sets, continua, and manifolds are all "best-case scenarios" in one way or another; they all have a high degree of simplicity and uniformity, and they all exhibit relatively little pathological behavior. Thus, objections to theories involving continua and/or manifolds are not automatically arguments for discrete theory. However, it is logical to explore the simplest alternatives first, and for this reason it is worth noting some additional doubts involving continuum and/or manifold theories. While no experiment achievable in the foreseeable future is likely to approach the **Planck scale**, it is believed by many that absolute physical limitations (such as the gravitational collapse of measuring devices) on experimentation may intervene as this scale is approached.[2] Even if this is not the case, it is still likely that nondiscrete theories will always hypothesize the existence of experimentally inaccessible entities. Of course, the existence of additional layers of structure is always possible whether they are experimentally accessible or not, but it seems better to add additional structure only if necessary rather than assuming an infinite regression of scales at which nothing new happens, or for which unfalsifiable hypotheses are made. Finally, **quantum nonlocality** can be viewed as direct evidence of the nonmanifold structure of spacetime, from the proper perspective. More precisely, nonlocal interpretations of experiments such as tests of **Bell's theorem** are based on the assumption of a metric structure for spacetime distinct from the causal structure, in the sense that events are viewed as being causally local (i.e. as directly influencing one another) without being metrically local (i.e. "nearby in space"). However, if all metric structure is defined in terms of causal structure, this distinction disappears along with the manifold structure of spacetime. This is the viewpoint provided by the **causal metric hypothesis**, which I turn to next.

The Causal-Metric Hypothesis.

Causality refers, of course, to cause and effect. Ideally, there is only one type of relation among the elements of a discrete causal theory, and these relations are taken to be **causal relations**. This approach depends on the assumption that the essence of causality can be captured by a binary relation, but discrete causal theory goes further by proposing that *all metric properties*, such as distance and time, also arise from the same binary relation. In particular, the most general versions of discrete causal theory abstain from any independent notion of **spacetime locality**. In particular, this is the case in causal graph theory and causal set theory. This identification of causality with a binary relation, and the proposal that metric structure emerges from the same relation is called the **causal-metric hypothesis**. Besides

[2]There are also good reasons to doubt the significance of the Planck scale, but analysis these reasons will be postponed in favor of a more detailed discussion later in the notes.

having the advantages of simplicity and economy, the causal-metric hypothesis automatically resolves problems arising from the existence of distinct causal and metric structures, such as **nonlocality** of **entangled states** in quantum theory and **time-travel paradoxes**[3] in general relativity. As with the absence of internal structure of the fundamental elements of discrete causal theory, the causal-metric hypothesis is sometimes relaxed in practice; in causal dynamical triangulation, for instance, the fundamental triangles are taken to be related both spatially and causally. Thus, from an abstract perspective, causal dynamical triangulation requires partially-directed **colored graphs**; i.e. graphs involving distinct symmetric and non-symmetric binary relations, for its description.

The causal-metric hypothesis, if correct, greatly simplifies and clarifies theoretical physics. In particular, it is the purest possible version of **background independence**, a subject of much tension between advocates of loop quantum gravity and string theory. Background independence means that the entire structure of a theory is dynamical; there is no static physical embedding space in which the dynamical entities of the theory reside. In this sense, general relativity, and hence loop quantum gravity, are background-independent, while quantum mechanics, quantum field theory, and string theory are background dependent. In fact, background dependent theories generally possess *three* distinct types of structures: **background structures**, **dynamical structures**, and **causal structures**. For example, string theory includes a background space, dynamical entities such as p**-branes** with metric structures, and the causal structure. Background-independent theories, as commonly understood, eliminate the background structure, while possibly retaining some degree of distinct metric structure. The causal-metric hypothesis goes further, by identifying metric structures as emergent entities arising from the causal structure.

The causal-metric hypothesis does not settle any of the controversial issues involving causality itself, only those arising from viewing the causal and metric structures of the universe as distinct. For instance, the issues of **self-causation**, **uncaused events**, **first causes**, **terminal effects**, **infinite causal regression** and **infinite causal progression** are all untouched by the causal-metric hypothesis. Indeed, as stated above, the causal-metric hypothesis merely states that what are ordinarily viewed as metric properties, both spatial and temporal, along with the usual notion of cause and effect, are all properties of a single binary relation; it says nothing about the actual properties of this relation. For example, **closed time-like curves** and **backward time-travel** exist in universes for which the binary relation has cycles, while these phenomena are absent if the relation is acyclic. Either type of theory is consistent, and no paradoxes arise. Rather, one must simply decide which theory or theories best agree with observation. As for any theory, empirical evidence may or may not uniquely discriminate among alternative explanations of singular, non-reproducible phenomena (such as the **origin of the universe**). In practice, physicists often make assumptions about some of these issues. For instance, self-causation is widely, thought not universally, viewed as pathological. In fact, in causal set theory,

[3]Confusion about time travel arises from two sources: conflation of time and causality, and uncertainty about the nature of causality itself. The first source of confusion is completely eliminated by the causal-metric hypothesis, and along with it any actual paradox. As detailed below, the second source of confusion reduces to (ideally testable) questions about the actual properties of the binary relation representing the causal structure, such as whether or not it is acyclic and whether or not it has maximal and minimal elements.

acyclicity of the causal relation is taken as one of the axioms. In my development of discrete causal theory, I will also focus on the case in which self-causation is omitted, but not exclusively. The other issues mentioned above, which involve **causal boundedness**, (uncaused events, first causes, etc.), cannot be similarly avoided in any reasonable causal theory, because of quantum theoretical and cosmological considerations.

Probability.

The incorporation probability into discrete causal theory is motivated by the philosophy that the statistical nature of quantum theory is fundamental and not merely a result of theoretical or experimental limitations. This is the opinion of the vast majority of physicists today. Apparent experimental violation of **Bell's inequalities** is regarded by many as conclusive proof of this viewpoint, although the experiments that have actually been done to date are not ideal and there remain a number of less obvious alternative interpretations. I have already alluded to Bell's theorem above as background to the suggestion, motivated by the causal-metric hypothesis, that "quantum nonlocality" is merely a misunderstanding resulting from the implicit assumption that spacetime has distinct metric and causal structures, but this does not mean that all **nonclassical** quantum phenomena can be explained away by discrete causal theory. In particular, there are many reasons, both experimental and philosophical, for accepting the **uncertainty principle** as a fundamental property of nature.

Of course, there are differences of opinion about how this apparent **indeterminism**[4] should be interpreted. Most physicists in the last half of the twentieth century regarded this behavior as evidence of a single, fundamentally statistical universe (the **Copenhagen interpretation** of quantum theory). More recently, partly (though not entirely) because of the influence of string theory, many have come to prefer the **many-worlds interpretation**, in which observed phenomena are regarded as merely one branch of a **quantum multiverse**, which includes every possible alternative. In any case, beginning with a probabilistic theory obviously does not rule out the possibility of a deterministic one; namely the special case in which a given probability is unity. Statistical universes, multiverses, and deterministic universes are all possible objects of study in discrete causal theory, depending on the form of the dynamical law.

Prediction in Causal Theory.

Physical theories are ultimately judged by their predictive power. Generally, though not invariably, predictions of interest concern the **future**. In particular, the emphasis on **experimental**

[4]Some prefer to view quantum mechanics as a *deterministic but acausal* theory, in the sense that the evolution of the state is deterministic and only the results of measurements are uncertain. In this view, the causes are insufficient to specify the effect; hence the term *acausal.* When I refer to quantum mechanics as nondeterministic, I mean, for instance, that there is no way to make a particle return a particular pair of values for consecutive position and momentum measurements. Either usage is fine provided it is explained clearly and used consistently.

verification and **reproducibility** in the **scientific method** calls for predicting, and to a limited extent, determining, the future. A typical problem is to predict the evolution of a system based on knowledge of its **present state**. For technical reasons, it is necessary in discrete causal theory to define "states" to include some of the **past** behavior of the system as well; at present I will ignore the details of how this is done, although this issue plays a significant role in the theory of entropy and in dynamics more generally, as discussed below. In temporal language, the present state of a system means the entire spatial extent of the system at a moment in time. This depends, of course, on the reference frame. In general relativity or string theory, the present state of a system is a **spacelike hypersurface**. The past is the set of spacetime events from which a signal sent at the speed of light could have reached the system. Using the causal-metric hypothesis, the problem can be restated in the language of causal theory as follows: the present state of a system is a *collection of causally unrelated elements of the causal structure*. In causal graph theory or causal set theory, such a collection is called an **antichain**. In causal dynamical triangulation, such a collection is a set of simplices sharing a given value of the discrete time variable. The past of a system in causal graph theory or causal set theory is the set of all elements connected by a causal chain to an element of the "present" antichain. In causal dynamical triangulation, the past is the set of all simplices connected to a "present" simplex by a sequence of simplices with strictly increasing time values.

In discrete causal theory, then, predicting the future means specifying the properties of collections of elements by means of knowledge about elements causally influencing those collections. This is illustrated schematically in figure 2 below. The yellow region represents the entire causal structure. The blue-and-red region represents the system under consideration, and the thick curve dividing the blue region from the red region represents the present state of the system. The blue region represents the known portion of the system, and the red region represents the portion of the system one wants to predict. The **causal direction**, which is the same as the **arrow of time** given the causal-metric hypothesis, points up the page, which is the usual convention. The schematic nature of the diagram makes no hint of discreteness, and indeed, such a diagram applies equally well to nondiscrete causal theory.

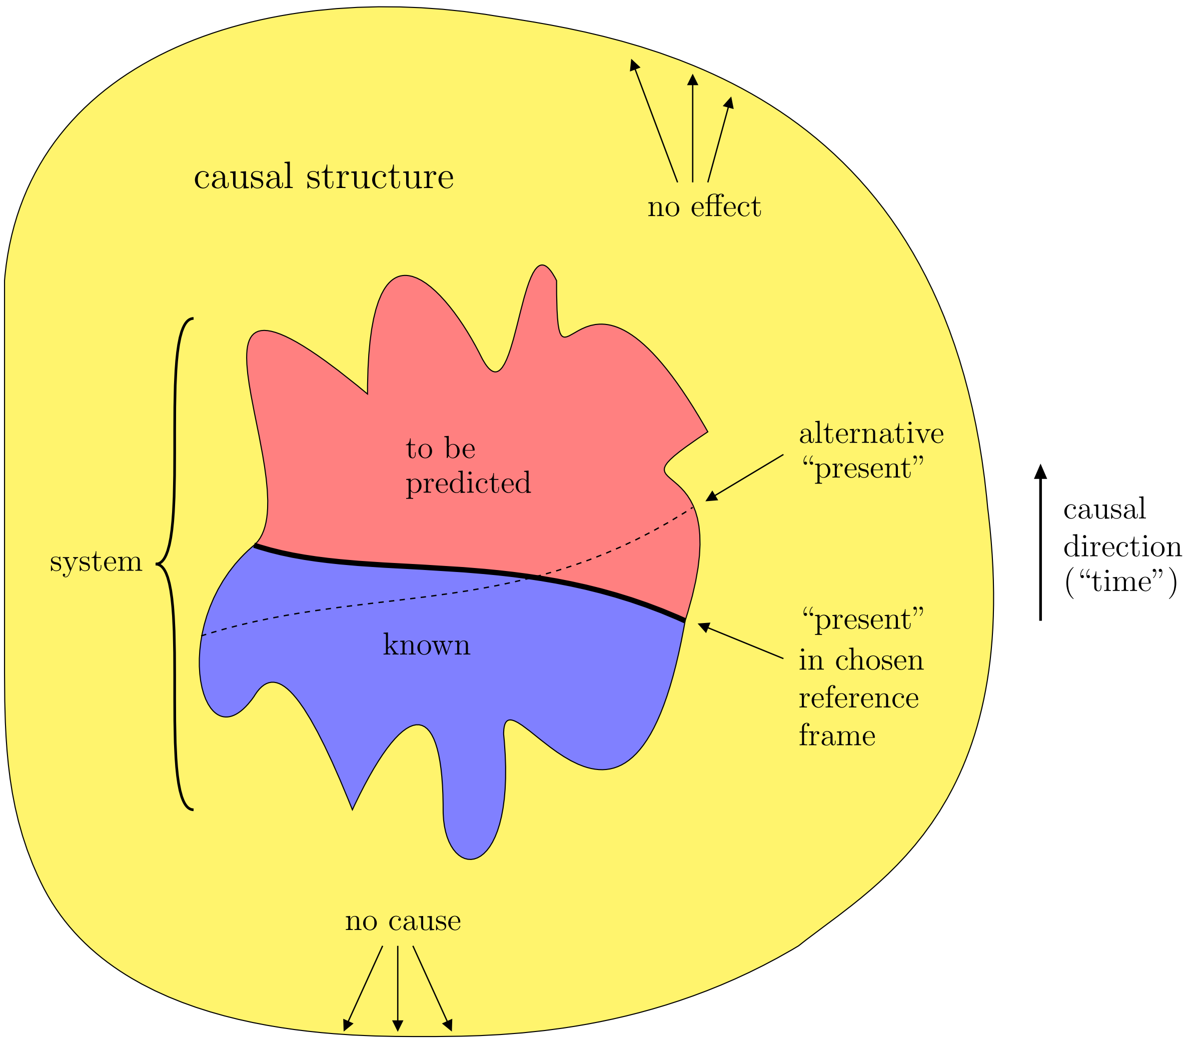

FIGURE 2

The notion of present depends on the point of view, or **reference frame**. The dashed curve in the diagram shows an alternative choice of present. Also, note that the diagram represents an entire history; the elements on the bottom boundary have no cause, and the elements on the top boundary have no effect. These boundaries may represent absolute **causal horizons** or may merely indicate the region of interest or accessibility. The region of interest in the present diagram is the red and blue area indicating the system under study. The lower boundary of the blue region does not indicate a causal horizon, but merely the limit of knowledge. Similarly, the upper boundary of the red region represents the "end of the experiment," not the end of time. For instance, it might represent the point at which data is no longer collected or analyzed.

The diagram represents a causal structure in an *ex post facto* manner, and this deserves some consideration. Even in general relativity, there has been some Zeno-like controversy[5] over whether

[5]See, for instance, *Relativity theory does not imply that the future already exists: a counterexample,* by Rafael Sorkin.

anything happens in the sense that viewing spacetime as a 4-dimensional manifold treats its existence as a fact rather than a process. It does not help to think of the manifold "developing over time," because time is part of its structure. I will explain the precise manner in which one may think of systems "developing" or "evolving" in discrete causal theory, but the above schematic diagram is not intended to indicate that anything in particular is determined about the future of the system, represented by the red region. Rather, the red region represents "whatever happens," which is what the observer with knowledge of the blue region wants to predict. I will now outline how this is done.

Discrete Causal Kinematics.

The first step is to describe, in very general terms, what is physically possible. This specification of possible behavior without referring to any dynamical law that might favor one causal structure over another is called **causal kinematics**. In causal kinematics, one begins with a particular causal structure and asks which larger causal structures the original structure can be embedded into in such a way that the "new elements" may be affected by, but may not affect, the elements of the original structure. Thus, the larger causal structures provide **possible futures** for the original structure. Figure 3, below, illustrates two possible futures for the known region in the previous diagram. The operation of adding a possible future is called a **pseudotransition**. Only two pseudotransitions are shown in figure 3, but there is usually an abundance, sometimes even an uncountable number, of pseudotransitions beginning from a given causal structure.

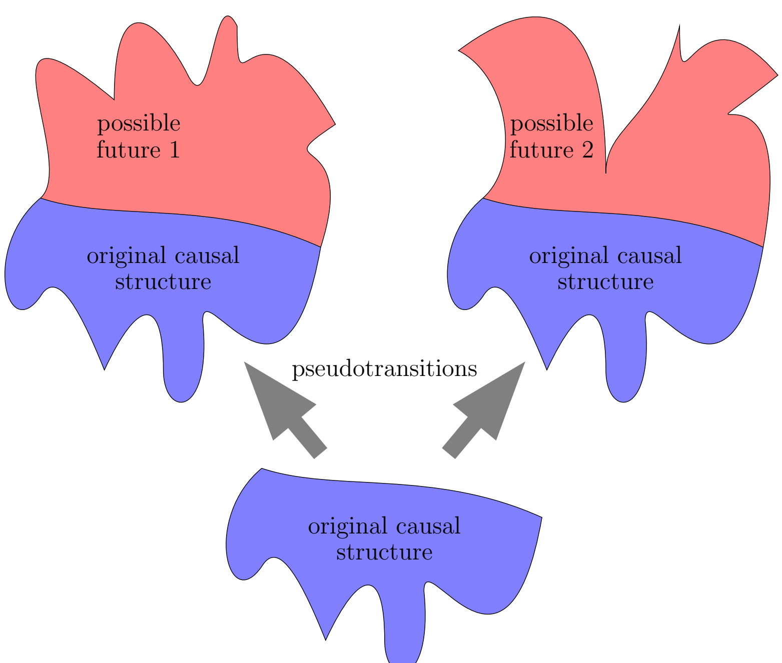

FIGURE 3

In discrete causal theory, it makes sense to talk about the **immediate future**. For example, in the causal graph in figure 1 on page 1, u is in the immediate future of x, but y and z are not in the immediate future of x. A **transition**[6] in discrete causal theory is a pseudotransition that adds only an immediate future to the causal structure. A composition of transitions is a pseudotransition, and is in fact a transition if the second transition adds no new elements causally related to the elements added by the first transition. In causal graph theory or causal set theory, a transition adds an **antichain**: a collection of elements with no causal relations among them. In causal dynamical triangulation, a transition adds a single "layer" of simplices. A given transition is either **reducible** into a composition of "simpler" transitions, or **irreducible.** Irreducible transitions are special because every transition, and hence every pseudotransition, can be written in terms of irreducible transitions. In causal graph theory or causal set theory, an irreducible transition adds a single element. In causal dynamical triangulation, an irreducible transition adds a single simplex.

In figure 4 below, the circles represent causal structures, and the line segments connecting them represent irreducible transitions. The numbers are labels, added for the purpose of distinguishing among causal structures, but have no physical meaning.

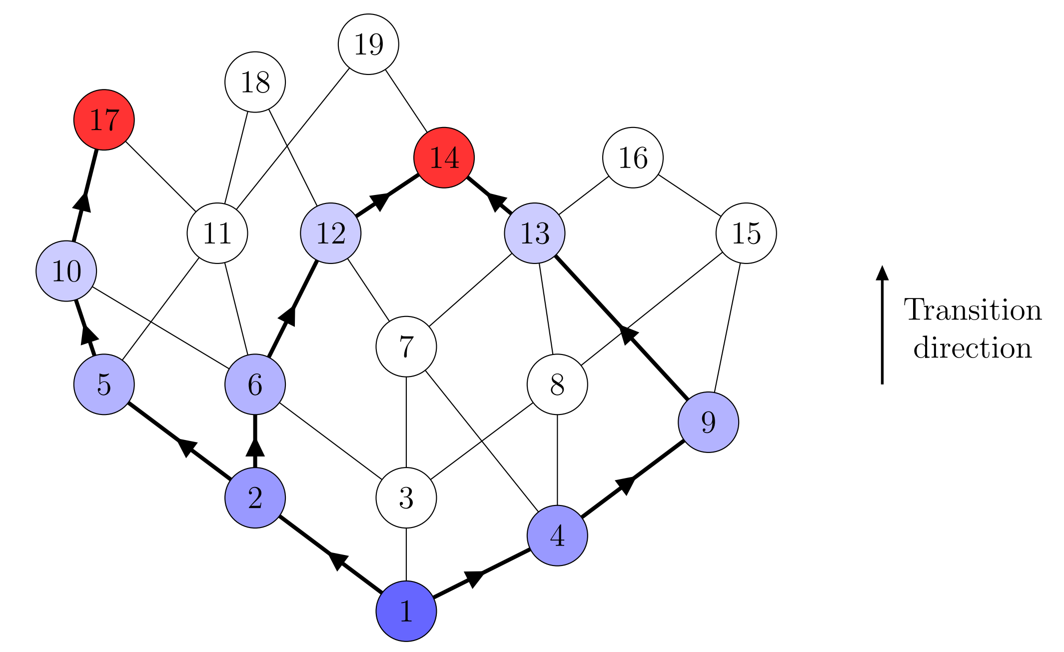

FIGURE 4

The direction of the transitions is up the page, so each causal structure strictly contains all structures connected to it by an ascending chain of line segments. For example, structure 14 strictly contains structure 13, which strictly contains structure 9, and so on. For a given discrete causal theory, the analogous diagram containing all possible causal structures and all irreducible transitions is called the **universal kinematic scheme** for the theory. It is "universal" because

[6]In causal set theory, the term transition is used to mean a pseudotransition that adds only a *single element*. This is what I call an irreducible transition; see below.

all irreducible transitions appear in the scheme. There are other useful kinematic schemes which are nonuniversal, but I will focus on universal kinematic schemes for the moment.

Figure 4 shows only a small part of a universal kinematic scheme, since universal kinematic schemes are necessarily infinite, as explained below. In very general terms, a universal kinematic scheme provides an "explanation" or "explanations" for every possible causal structure in the sense that tracing backwards from a given causal structure "tells a story" about the "evolution" of that causal structure. Usually there are very many distinct stories in this sense, because there are usually very many distinct descending paths beginning from any given causal structure in a kinematic scheme. All these stories are equally valid. For example, the paths $1 \to 4 \to 9 \to 13 \to 14$ and $1 \to 2 \to 6 \to 12 \to 14$ tell two different stories about the evolution of structure 14. Note that these two stories are *not* physically distinct histories, since each causal structure (in particular, structure 14) contains its entire history within itself. Rather, they are different points of view about the same history, analogous to reference frames in relativity. Looking back at the schematic diagram of the causal structure in figure 2 on page 4, one can see how families of "horizontal sections" represent reference frames in this sense. If one slices the causal structure into horizontal sections in two different ways, one obtains two different but physically equivalent methods of reconstructing the causal structure. This is illustrated in figure 5 below.

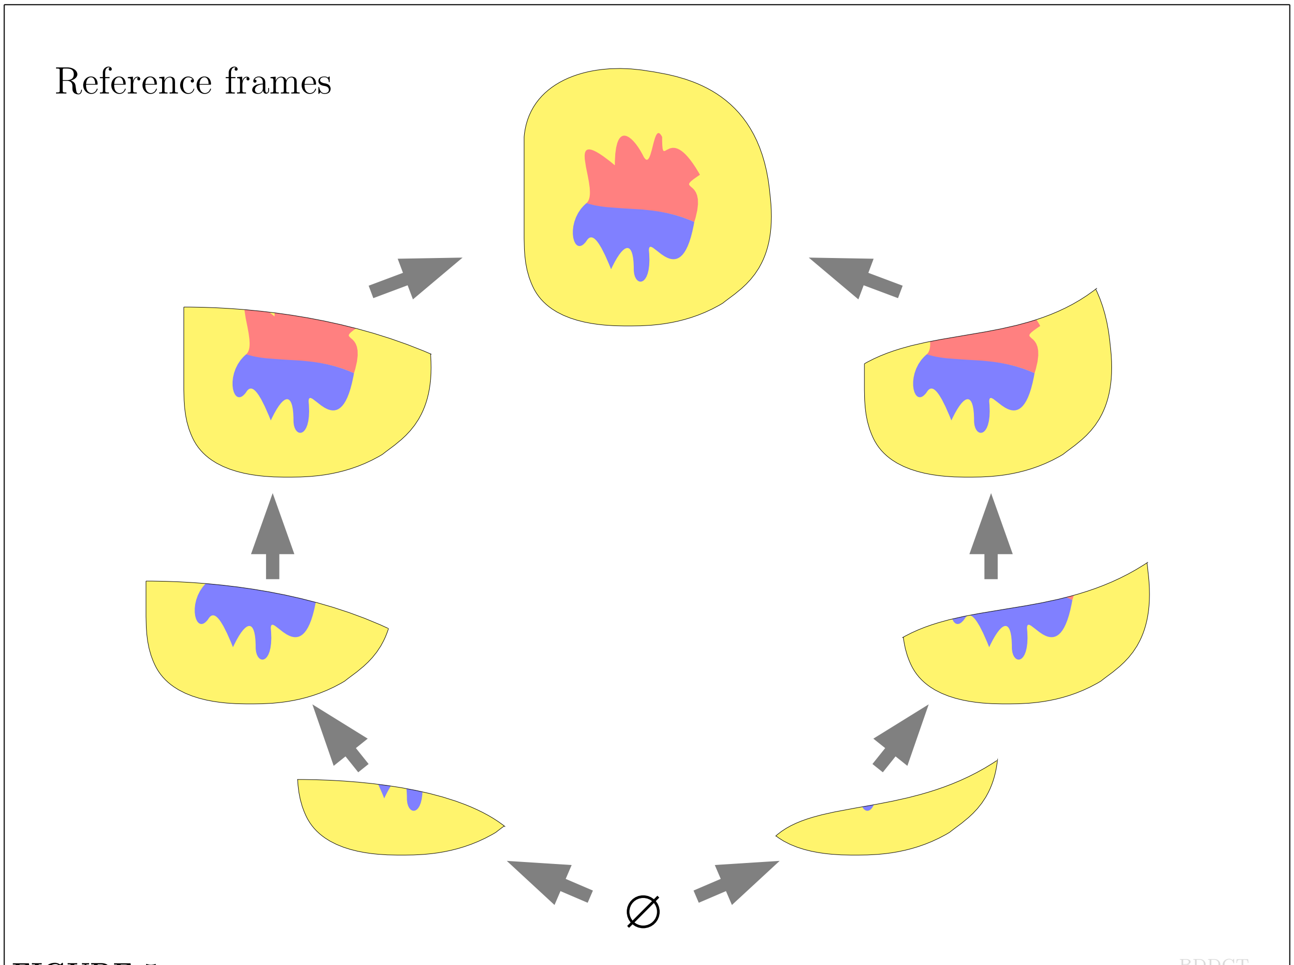

FIGURE 5

Ascending paths terminating at different causal structures in a universal kinematic scheme represent the evolution of *physically distinct* universes. For example, returning to figure 4, the paths $1 \to 2 \to 6 \to 12 \to 14$ and $1 \to 2 \to 5 \to 10 \to 17$ represent different physics, since the structures 14 and 17 are physically distinct. Thus, a universal kinematic scheme can be viewed as either a library of possible universes or a single **causal multiverse**. In multiverse interpretations of quantum mechanics, every statistical event is taken to represent a "branching off" of the universe into distinct histories, and this is similar to the idea of a universal kinematic scheme. An important difference is that different paths in a quantum multiverse always represent different physics (i.e. a quantum multiverse is a **tree**), while two paths sharing a common initial and terminal structure in a universal kinematic scheme represent identical physics. Another difference is that a quantum multiverse depends on an external time variable, while a universal kinematic scheme organizes objects with internal temporal (i.e. causal) structures.

You may have already noticed that a universal kinematic scheme itself has a "causal structure," where a causal structure B is in the **formal future** of a causal structure A if there is a sequence of irreducible transitions from A to B in the universal kinematic scheme. This multilevel structure built from causal structures of causal structures is one of the most interesting and useful features of causal theory.

Universal kinematic schemes are generally very large. In particular, given any pair of causal structures A and B, their disjoint union $A \sqcup B$ is again a causal structure, which is related to both A and B by pseudotransitions, and hence by compositions of irreducible transitions. Thus, there are paths from A and B to $A \sqcup B$ in the universal kinematic scheme. Therefore, no universal kinematic scheme has a maximal element, and every pair of elements in a universal kinematic scheme has a supremum. The empty causal structure is an infimum for any pair of causal structures, so a universal kinematic scheme is an **order lattice.**

Irreducible transitions generally do not give a very natural picture of the evolution of a discrete causal structure, since they add only "one element at a time" to the causal structure. It is physically more natural to think of entire **generations** of elements "appearing simultaneously," like spacelike hypersurfaces in relativity. For this reason, it is convenient to introduce other kinematic schemes which directly relate causal structures differing by a single generation, rather than by a single element. In general, a **kinematic scheme** is a directed graph whose vertex set consists of all possible causal structures and whose directed edges are transitions. There are some additional technical requirements on the edges which I will not discuss further at this point.

Discrete Causal Dynamics.

The final step in constructing a discrete causal theory is to specify a dynamical law from which to calculate which irreducible transitions are favored in the universal kinematic scheme. A **dynamical law** is a function μ whose source is the set of irreducible transitions in the universal kinematic scheme and whose target is a multiplicatively closed subset of a ring. Typical targets for a dynamical law μ are the **unit interval** $[0, 1]$ in $\mathbb{Q}$ or $\mathbb{R}$, or the **unit circle** S^1 in $\mathbb{C}$. In the first case, μ is generally a family of **probability measures**, one for the set of irreducible transitions from each causal structure. In the second case, μ is generally taken to specify a **phase factor** for each irreducible transition. Other choices are possible, but the ultimate purpose of specifying a dynamical law is to allow for (generally statistical) physical predictions, and the choices mentioned above are some of the more obvious possibilities pointing in this direction. For example, a dynamical law to the unit interval associates a probability with every path in the universal kinematic scheme, by multiplying together the images under the dynamical law of the irreducible transitions in the path. This produces a dynamics analogous to that of a **classical stochastic process**. By contrast, a dynamical law to the unit circle associates a phase factor with each path by multiplying the phase factors associated to each irreducible transition in the path. Interference of paths then produces a theory analogous to Feynman's **path integral formulation** of quantum theory. As discussed above, paths in the universal kinematic scheme with the same initial and terminal structures represent the same physics from different perspectives, so the concept of a sum over histories between two points in state space, where each path represents a different physical history of the system in question, must be modified to account for this extra information. In particular, recovery of the usual state space picture requires appropriate definitions of **states** and **state functions**, which are somewhat delicate.

Figure 6, below, shows a dynamical law to the unit interval on a portion of the universal kinematic scheme illustrated in figure 4. The probabilities assigned to the transitions from a given causal structure in the diagram always sum to unity because the dynamical law gives individual probability measures on each such set of transitions. More generally, there may be an infinite or even uncountable number of transitions from a given causal structure, so some care must be taken in properly analyzing sums of probabilities.

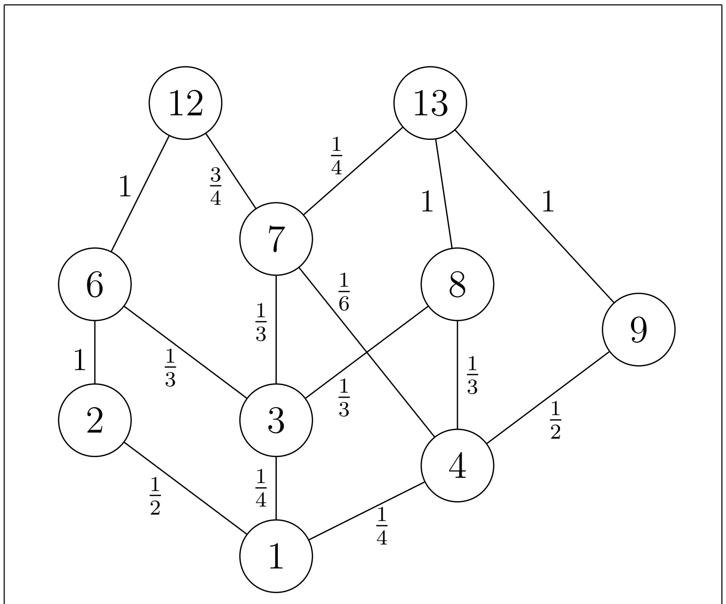

FIGURE 6

From the definition, it is clear that dynamical laws are deterministic only in special cases. One might expect determinism in a classical theory and statistical behavior in a quantum theory, but of course not every statistical theory is a suitable quantum theory. There is much to be said in this context about the essential distinctions between classical and quantum theories, the quantization of classical theories, and so on, but I will not belabor these issues here in the introduction.

It is a worthwhile exercise to consider carefully the physical meaning of the values associated to the irreducible transitions by the dynamical law in figure 6. For example, if one knows that a causal structure C is related to structure 12 by a pseudotransition, then it is physically meaningless to ask for the probability that the evolution of C included the transition $7 \rightarrow 12$ instead of the transition $6 \rightarrow 12$. This is because all kinematic stories explaining the evolution of C are equally valid. However, if one knows that a C is related to structure 7 by a pseudotransition, it makes sense to ask for the probability that C is also related to 12 by a pseudotransition. In this case, the probability is at least $\frac{3}{4}$. The probability may be greater than $\frac{3}{4}$; in particular, the universal kinematic scheme includes the disjoint union of structure 12 and structure 13, which is related to structure 13 by a pseudotransition, providing another possible way for structure 7 to "evolve" into a structure related to structure 12 by a pseudotransition. This is the sense in which the dynamical law allows for predicting the evolution of a causal structure. In general, the physically relevant information a dynamical law provides about a pair of causal structures A and B is the sum of the probabilities along all paths in the universal kinematic scheme beginning at A and intersecting some causal structure C related to B by a pseudotransition. Tracing back down from C through B to A, one is justified in regarding B as a stage in the evolution of A.

To do physics in a nonuniversal kinematic scheme, it is necessary to translate the values associated by the dynamical law to irreducible transitions in the universal kinematic scheme to appropriate values for transitions in the nonuniversal kinematic scheme. Alternatively, one could attempt to define a dynamical law directly on a nonuniversal kinematic scheme, but this requires some care. A nonuniversal kinematic scheme represents an arbitrary choice of viewpoint, and such a choice should not affect physical predictions. The requirement that the physical predictions of a dynamical law should not depend on arbitrary nonphysical information is called **covariance.** In a universal kinematic scheme, different reference frames for a given causal structure C are represented by different paths terminating at C. Since the values provided by a dynamical law depend only on individual irreducible transitions from C, without reference to paths terminating at C, covariance *a priori* imposes no conditions on the values assigned to irreducible transitions in a universal kinematic scheme.[7] However, covariance does impose conditions on how the representations of a dynamical law transform between nonuniversal kinematic schemes.

[7] Practitioners of causal set theory *do* impose conditions on probabilities in the universal kinematic scheme (the **positive sequential scheme** in this case) in the name of a principle referred to as **discrete general covariance.** I do not think this is a valid principle. More precisely, my view is that a principle similar to discrete general covariance applies at the level of the **tree of originary chains** on the universal kinematic scheme, in which each element encodes a causal structure *together with* an arbitrary choice of reference frame.

Methods For Deriving Dynamical Laws.

How does one go about determining an appropriate dynamical law? The ultimate judge is experimental evidence, but as for any theory involving extremely small scales, the process also involves the application of philosophical and mathematical principles, inspiration, and luck. I will briefly mention some of the more important principles here. A method which has been very useful in quantum theory, general relativity, and a number of more advanced theories is to first define a **classical theory**, often by means of an **action principle**, then **quantize** to obtain a **quantum theory**. The action principle is usually very simple conceptually, and intimately related to the basic structure involved. This provides hope of guessing the proper action directly. For example, the **Einstein-Hilbert action**[8] in general relativity depends only on the **curvature scalar** and the **metric tensor**. There is a discrete version of the Einstein-Hilbert action, called the Regge action, which is used in causal dynamical triangulation. There have also been efforts to modify the Einstein-Hilbert action for use in causal set theory. Despite its previous success, this approach suffers from some apparent limitations, the details of which are better left for a more complete discussion in a later section. In particular, there is difficulty in dealing with **topology change**, which is almost certain to play a role in any sufficiently general discrete causal theory. The relevant topology here is not, of course, the discrete topology, but the topology emerging from the causal structure.

Classical fields and **potentials** are among the basic ingredients of the above approach, and analogues of these concepts can be studied on their own merits without applying the entire machinery of the action principle and canonical quantization. Although fields and potentials are sometimes imposed as auxiliary structures, ideally they would arise from the causal structure itself. Otherwise, the binary relation representing the causal structure would be only part of the theory, and causal interpretations of interactions among the auxiliary fields would result in "multiple types of causality," which would be a depressing step backward from the simplicity of the causal-metric hypothesis. Fortunately, one can define **natural classical potentials** on causal structures whose forms are completely determined by the causal relation; indeed, taken together, they are equivalent to the causal relation. Since the formal structure of the classical dynamical laws of successful field theories such as Maxwell theory, Yang-Mills theory, and general relativity are similar to one another and are easily expressed in terms of the corresponding **vector potentials**, one might hope that a classical dynamical law could be written for discrete causal theory in terms of the natural classical potentials. However, since the dynamical laws alluded to above are differential equations whose form depends on the dimension of the spacetime manifold, the analgous equations for discrete causal theory are not easy to write down. In particular, the **dimension** and **curvature** of discrete causal structures are delicate subjects, especially locally. Since the discrete analogues of differential equations involve local properties, and since dimension and curvature play an indispensable role in successful field theories for manifolds, there is still some mystery involved in how to successfully complete this approach.

[8] Many view the Einstein-Hilbert action as only a low-energy approximation.

A second method by which one might try to deduce the dynamical law is by applying certain general physical principles that are believed on the basis of experience to be universal. Two such types of principles are **conservation principles** and **entropic principles**. Deciding what should be conserved is based partly on aesthetic grounds and partly on hypotheses about how familiar physical quantities (such as **angular momentum**) arise from the causal structure. One of the favorable aspects of discrete causal theory is that there are natural definitions of entropy for discrete causal structures, governed by simple counting arguments. Successful application of these definitions relies on appropriate definitions of states and state functions. This is not as simple as it might appear, since classical definitions of entropy rely on quantities involving time derivatives, whose discrete causal analogues involve the immediate past of a system as well as the present. Because of this, one obtains very different answers from different assumptions about the causal relation. In particular, the axioms of causal set theory lead to different result than my causal graph theory axioms. It seems possible that a generalized version of the **second law of thermodynamics** based on discrete causal entropy might be quite restrictive, and it is not out of the realm of possibility that such a law will eventually prove to be the *only* law necessary. The close relationship between gravity and thermodynamics has already been noticed from multiple points of view, and attempts to describe gravity by means of entropic principles are known collectively as **thermogravity**. Identifying an appropriate version of the second law of thermodynamics as the unique dynamical law of discrete causal theory would be the best possible scenario, since the resulting theory would have precisely one type of fundamental object, one type of interaction, and one law. Ideally, in light of the likely difficulty of obtaining direct experimental evidence, one would wish to link such a law to other conservation laws and derive it from multiple independent points of view (for instance, by means of the natural classical potentials mentioned above). Such efforts are ongoing at present.

Summary.

1. **Discrete causal theory** is an attempt to unify theoretical physics by replacing spacetime with a **discrete causal structure.**
2. In discrete causal theory, causality is fundamental, and spacetime is viewed as an emergent structure giving accurate approximations only at large scales.
3. Discrete causal theory has the advantages of conceptual simplicity, **background independence**, and **discreteness**.
4. To date, discrete causal theory is not well-developed.
5. Three important examples of discrete causal theory are **causal graph theory**, **causal set theory**, and **causal dynamical triangulation**.
6. The ingredients of discrete causal theory are **discrete sets**, **causality**, and **probability**.
7. Discreteness is a desirable property for a variety of reasons, including the success of discreteness hypotheses in quantum theory and the potential of discrete theories to avoid certain technical problems with existing nondiscrete theories.

8. Discrete causal theory is based on the **causal-metric hypothesis**, which states that causality and the metric structure of the universe are both manifestations of a single **binary relation**.
9. The incorporation of probability into discrete causal theory is based on the philosophy that the statistical properties of quantum theory represent fundamental principles of nature.
10. Prediction in discrete causal theory means specifying information about the properties of collections of elements of the causal structure by means of knowledge about elements causally influencing those collections.
11. **Discrete causal kinematics** is the description of what causal structures are physically possible, without referring to any dynamical law that might favor one structure over another.
12. Discrete causal kinematics involves specifying **possible futures** for every possible causal structure.
13. The operation of adding a possible future to a causal structure is called a **pseudotransition**.
14. In discrete causal theory, it makes sense to talk about the **immediate future.**
15. A **transition** in discrete causal theory is a pseudotransition that adds only an immediate possible future to the causal structure.
16. **Irreducible transitions** play an important role in discrete causal theory.
17. The directed graph whose vertices are all possible causal structures and whose edges are all irreducible transitions for a given causal theory is called the **universal kinematic scheme** for the theory.
18. Two directed paths with the same initial and terminal structures in the universal kinematic scheme represent the same physics, but from different **reference frames**.
19. The universal kinematic scheme may be viewed either as a library of possible universes or as a single **causal multiverse**.
20. The universal kinematic scheme itself has a causal structure given by the irreducible transitions.
21. Nonuniversal **kinematic schemes** often give a more natural physical picture than the universal kinematic scheme.
22. Nonuniversal kinematic schemes allow for transitions involving entire generations of elements, instead of one element at a time.
23. A **dynamical law** is a function from the set of irreducible transitions in the universal kinematic scheme to a multiplicatively closed subset of a ring.

24. Typical targets for a dynamical law are the **unit interval** $[0, 1]$ in $\mathbb{R}$ or $\mathbb{Q}$ and the **unit circle** S^1 in $\mathbb{C}$.
25. The purpose of a dynamical law is to give probabilities of various histories. Depending on the dynamical law, the resulting theory can resemble either a **classical stochastic theory** or a **quantum sum over histories**.
26. A standard method of deriving a dynamical law for a field theory is to begin with a **classical theory**, often defined by means of an **action principle**, then **quantize** the theory to obtain a **quantum theory**. This method may not work for discrete causal theory because of technical issues such as **topology change**.
27. Discrete causal structures have **natural classical potentials** analogous to the vector potentials in familiar field theories such as Maxwell theory, Yang-Mills theory, and general relativity.
28. A possible method of deriving a **dynamical law** for discrete causal theory is to work with the natural potentials directly. However, there are difficulties with this approach involving the analogues of manifold structures such as **dimension** and **curvature**.
29. Another method for trying to find an appropriate dynamical law is to apply general physical principles such as **conservation principles** or **entropic principles**.
30. Discrete causal structures have natural notions of entropy whose precise form depends on the definitions of **states** and **state functions**.
31. An appropriate version of the **second law of thermodynamics** for discrete causal structures may prove to be the desired dynamical law.