symmetry

IMPACT FACTOR 2.2

CITESCORE 5.4

Article

Counterintuitive Scenarios in Discrete Gravity Without Quantum Effects or Causality Violations

Benjamin F. Dribus, Jesiah T. Darnell and Neil Goldsmith

Special Issue

Symmetries in Gravity Research: Classical and Quantum

Edited by

Prof. Dr. Carlos Frajuca and Prof. Dr. Nadja Simão Magalhães

MDPI

https://doi.org/10.3390/sym16121609

symmetry

MDPI

Article

Counterintuitive Scenarios in Discrete Gravity Without Quantum Effects or Causality Violations

Benjamin F. Dribus [1,2,3,*], Jesiah T. Darnell [1] and Neil Goldsmith [4]

[1] Department of Mathematics, William Carey University, Hattiesburg, MS 39401, USA; jdarnell596471@student.wmcarey.edu
[2] College of Osteopathic Medicine, William Carey University, Hattiesburg, MS 39401, USA
[3] Foundational Questions Institute (FQxI), New York, NY 10163, USA
[4] Department of Chemistry and Physics, William Carey University, Hattiesburg, MS 39401, USA; ngoldsmith@wmcarey.edu
* Correspondence: bdribus@wmcarey.edu or bdribus@gmail.com

Abstract: In certain established approaches to quantum gravity, such as causal set theory and causal dynamical triangulations, discrete spacetime structure is taken to be a primary feature, not a secondary effect of "quantizing" a pre-existing classical continuum-based theory, as is done in approaches such as string theory and loop quantum gravity. For a priori discrete models, the full quantum theory is often obtained via some version of Feynman's sum-over-histories approach, in which each "history" is a discrete object viewed as a classical spacetime. Counterintuitive physical scenarios such as Schrödinger's cat or the grandfather paradox are typically associated with either quantum effects or causality violations, but we demonstrate that equally bizarre scenarios can arise at a purely classical level in the discrete causal context due to symmetry considerations. In particular, the graph-theoretic phenomenon of pseudosimilarity leads to situations in which alternative events occurring at physically distinguishable locations in the universe can cause different parts of the universe to "swap identities" in a fugue-like manner alien to continuum-based theories. This phenomenon is perhaps best understood as an extension of the relativity principle, which we call relativity of identity.

Keywords: quantum gravity; discrete spacetime; relativity; symmetry; pseudosimilarity; causal sets

check for updates

Citation: Dribus, B.F.; Darnell, J.T.; Goldsmith, N. Counterintuitive Scenarios in Discrete Gravity Without Quantum Effects or Causality Violations. *Symmetry* **2024**, *16*, 1609. https://doi.org/10.3390/sym16121609

Academic Editors: Carlos Frajuca and Nadja Simão Magalhães

Received: 29 October 2024
Revised: 27 November 2024
Accepted: 28 November 2024
Published: 4 December 2024

1. Introduction

A chief feature of modern physics is the prominence of counterintuitive scenarios arising from the consideration of scales, energies, relative velocities, and precisions far removed from ordinary experience. Relativity and quantum theory, the two main conceptual frameworks on which modern physics is based, are deeply colored by the influence of such scenarios, often in the form of specially-named "paradoxes". Among such scenarios sufficiently well known to have achieved broad popular awareness are the twin paradox, the barn-ladder paradox, Schrödinger's cat, and the grandfather paradox. None of these scenarios necessarily lead to paradoxes in a *logical* sense, although some of them challenge common philosophical presuppositions about what should be possible in a *physical* context. From a broader viewpoint, any new conceptual framework may be expected to engender scenarios of a type that would have escaped consideration in previous frameworks. Quantum gravity, which must necessarily combine the seemingly incompatible relativistic and quantum-theoretic frameworks, promises to live up to this expectation, even while the details of the theory remain uncertain. In the present study, we demonstrate that new types of strangeness indeed emerge from one particular conceptual building block of quantum gravity, whose consequences may be examined even without a complete theory, namely, discrete spacetime structure.

Roughly a century after relativity and quantum theory were originally established, many possible approaches to quantum gravity have been formulated, but only a few of these are broadly considered to possess strong physical motivation, largely due to the difficulty of obtaining definitive experimental evidence involving the necessary scales. The possibility of discrete spacetime structure plays a primary role in some of these theories and a secondary role in others. In theories emphasizing the fundamental causal relationships among individual events, such as causal set theory [1–3], causal dynamical triangulations [4–6], and several other theories [7–10], discrete spacetime structure is often taken to be a primary feature, while in theories more closely tied to pre-existing continuum-based methods, such as string theory [11] and loop quantum gravity [12], discrete structure is typically viewed as a secondary effect of "quantizing" a classical continuum theory. For a priori discrete models, a prominent method for obtaining a full quantum theory is to employ some version of Feynman's sum-over-histories approach [13], in which the continuum-based histories originally considered by Feynman are replaced by discrete objects such as causal sets or causal dynamical triangulations, each viewed as a classical history, and superposed in some way [14–18]. The new structural possibilities introduced by such models open the way to new types of unexpected behavior. These scenarios can arise from either primary or secondary sources of spacetime discreteness, but are most easily described in the a priori discrete context. Hence, to avoid unnecessary complications, we focus on simple discrete causal models called causal graphs, which are essentially equivalent to the "acyclic directed sets" studied in detail in [10], and which include causal sets as a special case. In considering such models, we are not proposing a new theory of spacetime, but rather focusing attention on structural elements common to a broad spectrum of existing approaches that are already well established in the literature.

Counterintuitive scenarios arising in any physical theory, whether actual or hypothetical, may be categorized in different ways depending on their degree of unexpectedness and whether or not their implications seem problematic to the theory's validity beyond a certain regime. In fact, one of the main practical reasons to consider such "paradoxes" is to aid in the evaluation of a theory by reducing the general to the particular. To some extent, the term "counterintuitive" is relative; for example, basic and well-established relativistic effects such as time dilation and length contraction are second nature to modern physicists, but may still surprise the general population. On the other hand, Böhr's view that anyone who is not shocked by quantum theory has not understood it remains quite relevant, and violations of causality still remain taboo to the point that they are often a priori ruled out in new theories, e.g., causal set theory, despite being quite permissible in general relativity. Hence, we may partition counterintuitive scenarios into a relatively "tame" Class A consisting of those that merely transgress ordinary experience, and a comparatively "wild" Class B consisting of those that seriously challenge even the most modern presuppositions.

We briefly examine a few examples of this categorization as a preliminary to the new counterintuitive scenarios we wish to consider in our present study. A prototypical Class A scenario is the twin paradox, which is merely a basic consequence of the relativistic distinction between inertial and highly accelerated motion. In this scenario, one twin remains on Earth while the other embarks on a lengthy journey through space, returning many years later from the perspective of Earth's approximately inertial frame of reference. Because less proper time has elapsed in the traveling twin's highly accelerated frame, the traveling twin is found to be many years younger than the Earthbound twin upon reunion. Such behavior is quite straightforward and thoroughly established, e.g., via cosmic ray experiments, despite the absence of suitable propulsion to carry out a full-scale version of the scenario. Another Class A scenario is the barn-ladder paradox, which is merely a colorful illustration of relativistic length contraction and the relativity of simultaneity. In this scenario, a farmer runs at relativistic speed into a barn, carrying a ladder whose rest length exceeds the barn's size. The "paradox" is found in the juxtaposition of the argument that the entire ladder should enter the barn before striking the back wall because the ladder Lorentz-contracts in the barn's frame of reference, and the counterargument that the ladder should strike the back wall with most of it still hanging out the front door because the barn Lorentz-contracts in the ladder's frame of reference. This issue is readily resolved via the relativity of simultaneity: one order of events occurs in the barn's frame, while the opposite order occurs in the ladder's frame, and the two are perfectly consistent. This type of behavior, again, is well understood and thoroughly established.

Schrödinger's cat falls into the Class B of "wild" scenarios, not because it implies a breakdown in quantum theory, but because its interpretation and bounds of applicability remain controversial. The crux of this scenario involves the hypothetical transfer of quantum-theoretic superposition of states to the macroscopic level by conditioning a macroscopic outcome on the distinction between two superposed microscopic eigenstates. The details are too horrible to relate, but the main issue is that such superpositions are not actually observed macroscopically, as far as we know. This, in turn, provokes a variety of candidate explanations for why this should be so, including "collapse of the wave function", "many worlds", "decoherence", and so on. We remark that cats also play a figurative role in some of the counterintuitive scenarios in our present study, but remain strictly alive and vocal throughout the proceedings despite suffering occasional crises of identity. The general-relativistic grandfather paradox is another needlessly violent Class B scenario. Evoked less objectionably in the film *Back to the Future*, it involves an individual who travels back in time, intersects his pre-history, and prevents his own birth. The "paradox" here is the apparently self-contradictory closed chain of causes and effects: the grandfather causes the grandson, who eliminates the grandfather, thereby eliminating his own cause, hence eliminating his ability to eliminate the grandfather, hence allowing his own existence, hence allowing him to eliminate the grandfather, and so on. Standard "solutions" include consistency criteria permitting self-causation without self-elimination, branching scenarios in which multiple copies or higher-dimensional slices of the individuals interact, and quantum-gravitational effects that prevent such closed chains from occurring in the first place. Causality violations like the grandfather paradox are often considered to be more problematic than superposition-related issues such Schrödinger's cat, in the sense that they are seen as more likely to indicate an actual breakdown of the theory rather than merely a challenge of interpretation.

As these examples illustrate, genuinely wild Class B scenarios in modern physics often involve either quantum effects or causality violations or both. In our present study, however, we find such scenarios arising naturally in a setting that involves *neither* of these common sources of trouble. These new scenarios fall into class B in the sense that they represent significant and poorly-understood departures from ordinary physical presuppositions, and their potential to emerge out of almost any conceivable type of causal structure guarantees them at least an abstract role in macroscopic events, even in situations without any direct relationship to quantum gravity. While a thorough explanation of these scenarios requires the mathematical background developed below, we can initially characterize them qualitatively as a sort of relativity of the identities of different parts of evolving causal structures. They are related to the relativity of simultaneity, depend strongly on symmetry principles, and involve the notion of distinguishability, but cannot be "explained away" by any conventional combination of these ideas. One way in which these scenarios manifest themselves is by contradicting the seemingly self-evident statement that "qualitatively different events occurring in identical systems lead to distinguishable histories". However, the way in which these histories conspire to avoid distinguishability involves a fugue-like mixing of the identities of their structural components that is totally unexpected and has no obvious continuum-based precedent. As a source of macroscopic analogies to render potential fundamental-scale interactions more comprehensible, we consider hypothetical histories involving cats, dogs, and human observers, whereby these effects may be translated into familiar terms. Rather than merely analyzing routine relativistic alternatives such as the order in which the various players in this drama take certain actions, we find ourselves deeply entangled with issues such as identifying *which animal is which* and determining whether or not these identifications remain stable as the system evolves.

Following the present Introduction, Section 2 covers the theoretical aspects of discrete causal gravity necessary to understand the relativity of identity at a precise level. These notions are well established in the literature and are merely included to introduce notation and to render the discussion as self-contained as possible. Section 2.1 offers a brief overview of the background and motivations for the discrete causal approach; the items of principal interest are the results of Hawking [19] and Malament [20] which imply that discrete causal structure is sufficient to recover the geometry of general relativistic spacetime at large scales, as explained by Sorkin and his collaborators [1]. Section 2.2 develops the basic mathematical characteristics of these discrete spacetime models, which we refer to here as causal graphs; causal sets are included as a special case. The book [10] gives a more comprehensive description of these objects using the terminology of "acyclic directed sets". Section 2.3 describes natural relationships called morphisms between pairs of causal graphs, which formalize the basic physical notions of evolutionary processes and symmetries in discrete spacetime. Section 2.4 develops the aspects of causal graph dynamics necessary to describe the scenarios we wish to analyze. Section 3 presents a detailed study of the relativity of identity. Section 3.1 introduces the crucial graph-theoretic notion of pseudosimilarity in mathematical terms. Section 3.2 describes one notable physical consequence of pseudosimilarity: qualitatively different events occurring in identical systems can lead to *identical* histories. Section 3.3 introduces and analyzes the relativity of identity as a second physical consequence of pseudosimilarity. Section 3.4 explores macroscopic analogies to further clarify the qualitative characteristics of scenarios involving the relativity of identity. Section 3.5 contextualizes the relativity of identity as an extension of the broad qualitative relativity principle epitomized by familiar notions such as length contraction, time dilation, and relativity of simultaneity, and explains its significance for spacetime, physical processes, and observers. Section 4 examines generalizations of these ideas and offers further discussion and analysis. Section 4.1 presents Theorem 1, demonstrating that the relativity of identity can apply to arbitrarily complex objects. Section 4.2 offers a symmetry-based explanation for the relativity of identity. Section 4.3 concludes with a general discussion, including a brief elaboration of applications of our study to several established approaches to fundamental spacetime structure, an overview of quantum-theoretic implications, remarks about phenomenology, and topics for further study.

2. Discrete Causal Gravity

2.1. Background and Motivation

The conceptual clarity of Einstein's general theory of relativity, enhanced by its experimental confirmation over the last century, offers a convincing argument that gravitation is a manifestation of spacetime structure rather than a force in the naïve sense. Spacetime models in all subsequent theories are thereby automatically endowed with gravitational implications. The notion of causality plays a leading role in many such models. In relativity, a causal curve is a curve whose tangent vector is everywhere timelike or null, meaning that the numeric quantity given by applying the metric as a bilinear form to the tangent vector does not change sign as one moves along the curve. Causal curves possess a natural notion of direction determined by the direction of their tangent vectors, with the physical interpretation that an event x can influence another event y only if there exists a causal curve from x to y. The familiar light cones or null cones appearing in Minkowski spacetime diagrams in special relativity delineate the regions of spacetime that can be reached by causal curves beginning or terminating at particular events, thereby specifying the domain of influence of each event. By convention, time flows "up the page" in such diagrams, so information is taken to propagate vertically from past to future. This convention makes sense for any causal model of spacetime, i.e., any model lacking closed causal curves which loop back to affect their own beginnings as in the grandfather paradox. We use this "up the page" convention throughout most of what follows.

The warping of null cones near massive bodies in general relativity exemplifies how the geometric effects of gravitation determine the causal structure of spacetime. In the late 1970s, it was gradually noticed that in certain important cases, the converse is nearly true: one can almost recover the entire geometry of a general relativistic spacetime manifold simply from the knowledge of which events do or do not affect other events. These notions are made precise by the "metric recovery theorems" of Hawking [19] and Malament [20]. The missing data, which prevent complete metric recovery, involve scale, i.e., the relative sizes of different regions of spacetime. Around this time, several physicists [21–23], most notably Sorkin and his collaborators [1], realized that if spacetime were composed of individual grains or corpuscles of roughly the same size, then the necessary scale data would be available "for free" by simply counting the number of grains in any given region. Sorkin neatly encapsulated this realization via the statement, "Order plus number equals geometry". Following this idea, the entire theory may be described in purely causal terms by specifying which pairs of grains are causally linked. An extension of the same idea is the causal metric hypothesis [10,24], which states that the properties of the physical universe in general, and the metric properties of spacetime in particular, are manifestations of causal structure. In view of the metric recovery theorems, it is natural to apply this philosophy to spacetime exclusively, attributing particle theory to some additional structure of "fields" attached to grains of spacetime and/or their relations. This is how matter and energy are typically treated in causal set theory, for example. However, the potential irregularity of discrete causal structure, in contrast to the local uniformity of manifolds, offers the opportunity to try to model *all of physics* via unified causal structure, with particle properties emerging from specific types of local configurations. The optimistic assertion that this is possible is described as the "strong form" of the causal metric hypothesis in [10]. While the potential symmetry-breaking role of additional fields attached to a causal structure has some possible implications for the scenarios we consider here, we mostly ignore these complications in our present study and focus on causal structure itself.

The resulting spacetime models consist of nodes, representing the grains of spacetime, physically interpreted as individual events, connected together by directed edges representing causal influences between pairs of events. Sequences of such edges are analogous to causal curves. Mathematically, these models are special types of directed graphs, which generalize the causal sets initially proposed by Sorkin et al. As a refinement, it is common to rule out the existence of cycles, which are discrete causal analogues of closed causal curves that can lead to issues such as the grandfather paradox. Imposing this prohibition is not an open-and-shut matter, due to the prevalence of closed causal curves in general relativity, even in very basic and generic situations such as the Kerr metric [25]. However, closed causal curves are manifestly not the norm in ordinary physics, and it seems reasonable to set aside the complications of considering them at least on a tentative basis. We therefore assume unless stated otherwise that our models are devoid of such cycles. This leads us to restrict attention to acyclic directed graphs, which we shorten to "causal graphs" due to the physical interpretation of these objects in our present study.

Figure 1, modified from a similar figure in [10], illustrates how classical spacetime structure is viewed as a causal graph in discrete causal gravity. On the left, spacetime structure is depicted at relatively large scales, as it appears in the context of general relativity. At such scales, known via particle accelerators such as the Large Hadron Collider to include distances down to around 10^{-20} m or less, spacetime appears to be smooth. The warping shown in the figure represents the nontrivial curvature of spacetime in general relativity due to the presence of matter and energy. The curved grid lines represent a choice of coordinate system, which provides local delineations between "space" and "time". The central node, labeled x, represents a particular spacetime event, and the dark triangular or cone-shaped regions above and below x represent its past and future, respectively. The past of x consists of all spacetime events from which a causal curve can reach x. These are the events that can possibly influence x. The future of x consists of all spacetime events that may be reached from x by a causal curve. These are the events that x can possibly influence. The black curve is an example of a causal curve through spacetime near x; if a physical object moves along this curve, then the curve is called the object's world line.

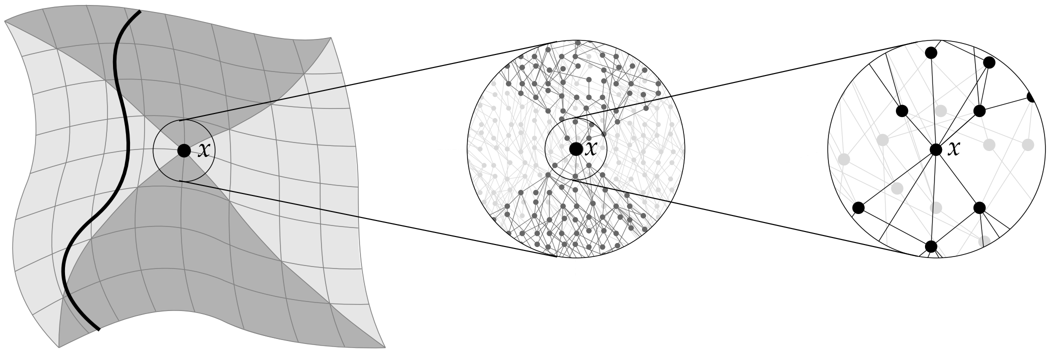

Figure 1. Discrete causal theory hypothesizes that spacetime is not infinitely subdivisible, but that granularity begins to manifest itself at small scales. At the fundamental scale, nothing remains but individual events and their causal relationships. In aggregate, they form a *causal graph*.

The middle portion of Figure 1 shows a "mesoscale" view of spacetime at which granularity begins to manifest itself. The existence of such granularity is one of the principal hypotheses of discrete causal theory. Optical experiments place significant constraints on what degrees and types of granularity can be consistent with observations, but many scenarios remain in which such granularity could be present [26]. Expectations vary regarding the precise scale at which this might occur; 10^{-35} m is a reasonable guess based on the rough heuristics offered by the Planck scale, which involves considerations such as the gravitational stability of minimal energy units. However, arguments have been raised for the possibility of departures from a continuum structure at somewhat larger scales [27,28]. In the mesoscale setting, the irregularity of the causal past and future of x begin to become apparent: rather than smooth cone-like regions of a manifold, they consist of complex networks of discrete events and relationships among them. The right-hand portion of the figure illustrates the hypothesized microscale structure of spacetime, in which individual causes and effects near x are clearly resolved. Black nodes represent events that *actually* influence x or are influenced by x, while gray nodes represent events unrelated to x, which may or may not be related to each other. A key point in understanding this hypothesized structure is that the nodes representing events and the edges representing causal relationships are not considered to exist *inside* any type of space or spacetime whatsoever; rather, "spacetime" is viewed as merely an approximate way of conceptualizing and describing the aggregate structure of vast numbers of such nodes and edges.

Many refinements to the basic notion of causal graphs are possible. The entire class is immense, and many of its members are manifestly unsuitable as models of the familiar four-dimensional spacetime suggested by observation. Special classes of such graphs used for particular spacetime models have been given different names; for example, Sorkin's causal sets involve a few additional assumptions such as transitivity and interval finiteness, while causal dynamical triangulations incorporate much more restrictive assumptions that essentially hardwire the correct dimensionality of the resulting spacetime at small scales. For our present purpose, we need not consider the details of such additional assumptions. We merely remark that the counterintuitive scenarios presented here can arise for most such models, and therefore apply to a broad spectrum of well-established approaches to fundamental spacetime structure. Furthermore, Theorem 1 below implies that such scenarios can involve arbitrarily complex objects, and do not simply represent odd pathologies. In what follows, the term "graph" will usually mean "causal graph" without specific additional assumptions, besides some elementary finiteness conditions.

2.2. Causal Graph Basics

We now give the basic definitions necessary to precisely understand such causal graph spacetime models, before examining some of the surprising structural possibilities that arise from this viewpoint, culminating in the relativity of identity. We emphasize that we are *not* introducing a new theory of spacetime here, but merely isolating some common structural elements of several well-established approaches, including causal set theory. We use the symbols G, G', H, H', etc. to denote graphs of various types, usually causal graphs. It does no harm to note that G may be considered to stand for "graph" and H may be considered to stand for "history". When considering two or more such graphs simultaneously, we may either use G and H, or we may use "prime notation" such as H, H', and H''. It is important not to ascribe any fixed physical meaning to specific symbols in this context, since a variety of different roles must be described. We begin by precisely defining directed graphs in general. We restrict ourselves to the case of finite graphs for the purposes of this study. Infinite graphs can be interesting as physical models, but are unnecessary for the scenarios we wish to explore. In particular, it is always possible to restrict attention to a local neighborhood which may be modeled via a finite graph under reasonable assumptions.

Definition 1. *A* **directed graph** *H consists of a finite set of nodes V, a finite set of edges E, and* **initial** *and* **terminal node maps** *$\alpha, \omega : E \to V$, such that for each pair of nodes $v, v' \in V$, there is at most one edge e satisfying $\alpha(e) = v$ and $\omega(e) = v'$. A* **subgraph** *G of H consists of subsets $U \subset V$ and $D \subset E$, with initial and terminal node maps given by restricting α and ω. G is called a* **full subgraph** *of H if D includes all the edges from H whose initial and terminal nodes belong to U. The* **complement** *of G is the full subgraph of H induced by the complement node set $V - U$; i.e., it includes all edges between pairs of nodes not in U.*

Figure 2 illustrates a simple-looking directed graph H, which we regard throughout our study as representing a discrete causal history, and whose properties will be seen to raise profound questions about fundamental spacetime structure. The figure demonstrates that the recipe given by Definition 1 is actually quite straightforward, despite its a priori abstract flavor. The initial and terminal node maps α and ω may be regarded as providing instructions for how to "snap together" the nodes and edges of H, akin to how labels guide concrete constructive projects such as building "ball and stick" models in elementary chemistry. We have labeled the nodes and edges in a suggestive way that makes the reading of the images of α and ω automatic: the initial node $\alpha(e_{ij})$ of each edge e_{ij} is just v_i, while the corresponding terminal node $\omega(e_{ij})$ is just v_j.

(A)

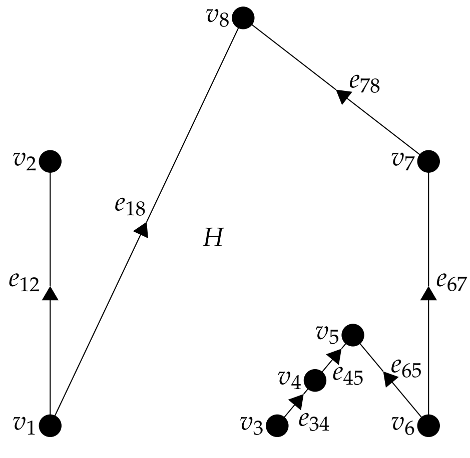

(B)

Instructions from initial and terminal node maps:

$\alpha(e_{12}) = v_1, \quad \omega(e_{12}) = v_2$
$\alpha(e_{18}) = v_1, \quad \omega(e_{18}) = v_8$
$\alpha(e_{34}) = v_3, \quad \omega(e_{34}) = v_4$
$\alpha(e_{45}) = v_4, \quad \omega(e_{45}) = v_5$
$\alpha(e_{65}) = v_6, \quad \omega(e_{65}) = v_5$
$\alpha(e_{67}) = v_6, \quad \omega(e_{67}) = v_7$
$\alpha(e_{78}) = v_7, \quad \omega(e_{78}) = v_8$

Figure 2. (**A**) Directed graph H, viewed as a discrete causal history; (**B**) initial and terminal node maps α and ω provide instructions for "snapping together the nodes and edges".

We pause briefly to spell out the physical interpretation of H. The nodes $v_1, v_2, \ldots, v_8$ represent individual events, which may be regarded in the familiar relativistic manner as "single points in space at single instants in time". At a fundamental level, of course, this association is reversed: "space" and "time" are regarded as emergent concepts describing the aggregate behavior of large numbers of such nodes, together with the edges connecting them. Considerations arising from the Planck scale, together with Sorkin's recipe for determining spacetime volume by counting events, lead us to ascribe a small nonzero size to each event. In spatial terms, for lack of definitive evidence, we may tentatively think of such events as possessing a Planck-scale size of roughly 10^{-35} m. For technical reasons, it is necessary to allow for some small fluctuations in this basic size [26]. Of course, this association is also reversed at the fundamental scale: "size" is regarded as an emergent concept describing the number of such events involved in a given process. Each edge e_{ij} in Figure 2 represents a fundamental directed causal relationship between the events represented by the nodes v_i and v_j. In other words, the event represented by v_i is partially responsible for the genesis of the event represented by v_j. For example, v_1 causes v_2, while the combination of v_1 and v_7 causes v_8, with indirect causation from v_6 transmitted via v_7. Since "time" is merely regarded as a way of describing which events cause other events, we say that v_i "precedes" v_j whenever there is an edge from v_i to v_j. Planck-scale considerations lead us to tentatively ascribe a time interval of roughly 10^{-43} s to each edge. As usual, this association is ultimately reversed: one second of proper time is expected to consist of approximately 10^{43} consecutive fundamental relations.

Since the directed graph in Figure 2, taken as a whole, is understood as a model of spacetime, it may be regarded as a very simple "universe" or "history". The term "universe" conveys the idea that the system is causally closed: all causes impacting the events represented by the nodes of the graph are themselves represented by the edges of the graph. However, this viewpoint ignores quantum effects, which are generally introduced by considering the superpositions of such graphs, following the general sum-over-histories approach to quantum theory pioneered by Feynman [13]. This explains the use of the term "history", or more precisely, "classical history", for such a graph, since quantum histories are also studied in various contexts. In the present setting, each graph, with a suitable notion of equivalence, is considered to represent a different classical history, while the entire quantum-theoretic universe is constructed from a configuration space of such histories via the superposition principle. Besides a brief discussion in Section 4.3, we mostly ignore quantum effects here; indeed, the surprising character of the scenarios we analyze partly arises from the fact that quantum effects are *not* necessary to produce highly counterintuitive outcomes. At a practical level, it is useful to apply to such models the intuition derived from Minkowski spacetime diagrams in relativity. For example, we can say that the three events represented by v_1, v_3, and v_6 all occur at the "beginning of time", since no events precede them, while the three events represented by v_2, v_5, and v_8 all occur at the "end of time", since no events succeed them.

Resuming our technical development, we sometimes find it useful to denote a directed graph H explicitly by writing its structural ingredients as a quadruple $H = (V, E, \alpha, \omega)$. A subgraph G of H with node set $U \subset V$ and edge set $D \subset E$ uses the same initial and terminal node maps as H, appropriately restricted, and may therefore be denoted by the quadruple $G = (U, D, \alpha, \omega)$. The use of the symbol V for the node set of H is a carryover from abstract graph theory, where nodes are often called vertices. The edge set E may be defined alternatively as a subset of the Cartesian product $V \times V$; this definition is equivalent to the "independent" notion of E appearing in Definition 1, due to the restriction that a given pair of nodes can serve as the initial and terminal nodes of at most one edge. We use the latter definition because it generalizes better and meshes with the literature.

We next define directed paths in directed graphs, which provide the discrete causal analogues of causal curves.

Definition 2. *A* **directed path** γ *from a node* v *to a node* v' *in a directed graph* H *is a sequence* $(e_1, e_2, \ldots, e_n)$ *of edges such that* $\alpha(e_1) = v$, $\omega(e_m) = \alpha(e_{m+1})$ *for each* $m \in \{1, 2, \ldots n-1\}$, *and* $\omega(e_n) = v'$. *The nodes* v *and* v' *are called the* **initial** *and* **terminal nodes** *of* γ*, respectively.*

The notion of initial and terminal nodes of a directed path extends the definition of initial and terminal nodes of individual edges, which are paths of length one. The reason for using just a single index for the edges in Definition 2 rather than the double-index notation used in Figure 2 is that, in the present case, we do not need to refer explicitly to the initial and terminal nodes of each edge. The sequence of edges (e_{67}, e_{78}) in Figure 2 is a directed path of length 2 from the node v_6 to the node v_8. In physical terms, the existence of this path is interpreted as representing the flow of influence or information from v_6 through v_7 to v_8. By contrast, the sequence of edges (e_{65}, e_{67}, e_{78}) is *not* a directed path from v_5 to v_8, because the order of the sequence does not agree with the directions of the edges. In physical terms, influence or information does *not* flow from v_5 to v_6, and thence to v_8, because the edge e_{65} points in the wrong direction. Using the double-index notation for edges from Figure 2, a given sequence of edges qualifies as a directed path if and only if the second index of each edge coincides with the first index of the next edge. The sequence of edges (e_{65}, e_{67}, e_{78}) does define an *undirected* path from v_5 to v_8, but the reasons to consider such paths in the present context are limited because they fail to respect the causal structure. One legitimate reason is that undirected paths determine topological connectivity, and connectivity properties can be useful in distinguishing different parts of graphs in some of the examples below. However, since directed paths generally dominate the discussion otherwise, we usually drop the term "directed" and assume that all paths under consideration are directed unless stated otherwise.

Having established the notion of directed paths as discrete causal analogues of causal curves, with the role of transmitting influence or information, we next turn to the problem of defining which regions of a directed graph H transmit or receive such influence or information to or from a specified subgraph. These regions are the discrete causal analogues of the "pasts" and "futures" of spacetime regions in general relativity.

Definition 3. *Let $H = (V, E, \alpha, \omega)$ be a directed graph and $G = (U, D, \alpha, \omega)$ a subgraph. The* **past** *G^- of G in H consists of the full subgraph induced by the set of all nodes $v \in V$ such that there exists a path beginning at v and terminating at some node $u \in U$. The* **future** *G^+ of G in H consists of the full subgraph induced by the set of all nodes $v \in V$ such that there exists a path beginning at some node $u \in U$ and terminating at v.*

In most cases, it is permissible to think about pasts and futures simply in terms of families of nodes, rather than invoking the entire structure of the full subgraphs induced by these families. For example, the past of the node v_4 in Figure 2 is the single node v_3, and its future is the single node v_5. The future of the node v_6, in terms of nodes, is just $\{v_5, v_7, v_8\}$; including the full subgraph structure adds the edges e_{65}, e_{67}, and e_{78} to the picture. In general relativity, there are distinctions between causal and temporal pasts and futures, as well as considerations involving the inclusion of boundaries, such as whether or not a given event should be considered to be part of its own past or future. These complications can generally be ignored in the present context.

We next define cycles in a directed graph, which are analogous to closed causal curves in general relativity. From a physical perspective, cycles may be interpreted to represent causality violations in which events affect their own pasts. As stated above, we choose to eliminate such cycles from consideration in our present study, although the precedent of general relativity cautions against disregarding them absolutely. The counterintuitive scenarios we examine here therefore involve *neither* quantum effects nor causality violations. Ruling out cycles completes our definition of causal graphs, interpreted as discrete causal histories.

Definition 4. *A* **cycle** *in a directed graph $H = (V, E, \alpha, \omega)$ is a directed path $(e_1, e_2, \ldots, e_n)$ such that $\alpha(e_1) = \omega(e_n)$. H is called an* **acyclic directed graph***, or* **causal graph***, if it has no cycles.*

To summarize, we refer to finite acyclic directed graphs as causal graphs in our present physical context, with the understanding that they are to be interpreted as models of classical spacetime without causality violations. Due to the acyclic property, such graphs can always be drawn in the same manner as Minkowski spacetime diagrams in special relativity, in which the arrow of time or direction of information flow is "up the page". We may therefore eliminate the explicit use of arrows on the edges, such as those appearing in Figure 2, since the causal relationship represented by each edge is always interpreted to flow from the "lower" node to the "upper" node. We follow this convention throughout the rest of the paper.

2.3. Causal Graph Morphisms and Symmetries

Just as spacetime manifold symmetries play a central role in special and general relativity and in particle theories such as the standard model based on relativistic spacetime structure, so causal graph symmetries play a central role in discrete causal physics. In technical terms, symmetries of causal graphs are certain types of "natural self-relationships" called automorphisms, which are a special case of natural relationships called morphisms between pairs of causal graphs. The way in which causal graph morphisms are "natural" is the same way in which linear transformations of vector spaces, homeomorphisms of topological spaces, diffeomorphisms of smooth manifolds, and homomorphisms of groups or rings are natural: they are maps between pairs of objects that respect the internal structure of the objects, in this case, initial and terminal nodes. Besides encoding symmetries, morphisms also serve as building blocks for describing evolutionary processes for discrete spacetime. Since the definition of morphisms does not depend on the absence of cycles, we present it in the general context of directed graphs.

Definition 5. *A* **morphism** $\phi : H \to H'$ *from a directed graph* $H = (V, E, \alpha, \omega)$ *to a directed graph* $H' = (V', E', \alpha', \omega')$ *is a pair of set maps* $\phi_V : V \to V'$ *and* $\phi_E : E \to E'$ *that respect initial and terminal node maps in the sense that* $\alpha'(\phi_E(e)) = \phi_V(\alpha(e))$ *and* $\omega'(\phi_E(e)) = \phi_V(\omega(e))$ *for every edge* $e \in E$*. The* **image** $\phi(H)$ *of* ϕ *is the subgraph of* H' *with node set* $\phi_V(V)$*, edge set* $\phi_E(E)$*, and initial and terminal node maps defined by restricting* α' *and* ω' *in the obvious way. The* **identity morphism** Id_H *is the pair of identity set maps* $(\mathrm{Id}_V, \mathrm{Id}_E)$*. A morphism* $\phi : H \to H'$ *is called an* **isomorphism** *if there exists an* **inverse** *morphism* $\psi : H' \to H$ *with the property that* $\psi \circ \phi = \mathrm{Id}_H$ *and* $\phi \circ \psi = \mathrm{Id}_{H'}$*. A morphism* $\phi : H \to H'$ *is called a* **monomorphism** *or* **embedding** *if the induced morphism from* H *to* $\phi(H)$ *is an isomorphism. An* **automorphism** *of a directed graph* H *is an isomorphism from* H *to itself.*

Whether or not two given directed graphs are isomorphic can be difficult to determine, regardless of how they are defined or described. The "graph isomorphism problem" is a famous problem in complexity theory that is known to belong to the class NP, and that has attracted much recent interest due to significant progress by Babai [29] in developing algorithmic methods for determining isomorphism. Figure 3 illustrates an explicit isomorphism ϕ between two causal graphs H and H', encoded by labeling the nodes and edges of the two graphs to indicate which structural components of H' are the images of the corresponding structural components of H. In terms of the double-index notation introduced in Figure 2, ϕ is defined by setting $\phi_V(v_i) = v_i'$ and $\phi_E(e_{jk}) = e_{jk}'$ for all appropriate indices i, j, k. The inverse morphism ψ is defined in the obvious way by setting $\phi_V(v_i') = v_i$ and $\phi_E(e_{jk}') = e_{jk}$. Note that the isomorphism ϕ is "non-obvious" in the sense that the abstract structure of H and H' does not appear to be "the same" at first glance, and the structural identifications made by ϕ are not what one would naïvely expect in some cases. For example, the node v_5 may appear naïvely to play essentially the same role in H as the node v_8' does in H', but its actual structural role is identical to that of v_5'. These details are important for what follows.

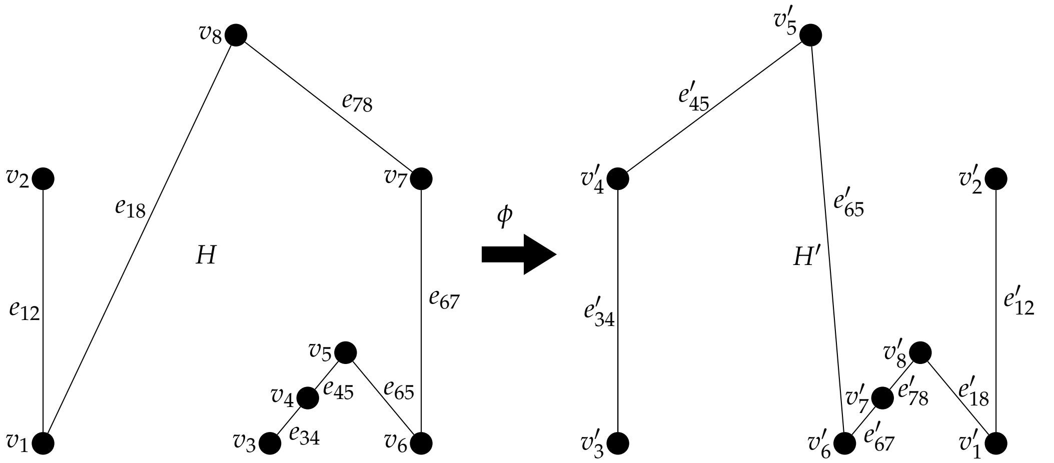

Figure 3. Isomorphism ϕ between two causal graphs H and H', encoded in terms of labels; for example, $\phi_V(v_1) = v'_1$, and $\phi_E(e_{12}) = e'_{12}$.

A useful intuitive way to think about graph isomorphism is in terms of stretching, shifting, and/or reflecting parts of a graph model, which one imagines to be constructed of flexible material such as rubber bands. If a rubber band model of one graph H can be stretched, shifted, and/or reflected so that it matches the shape of another graph H' while preserving the directions of the edges, then the two graphs are isomorphic. In terms of the "up the page convention" for causal graphs inherited from Minkowski spacetime diagrams, the rules for such manipulations are simply that the rubber bands cannot be broken and that each pair of nodes must maintain their original vertical order throughout the process. For example, Figure 4 visually demonstrates that the graphs H and H' from Figure 3 are indeed isomorphic via a sequence of such manipulations. Isomorphism means that H and H' are considered to be *identical for physical purposes* in the context of discrete causal gravity, since they represent the same family of causal relationships among events.

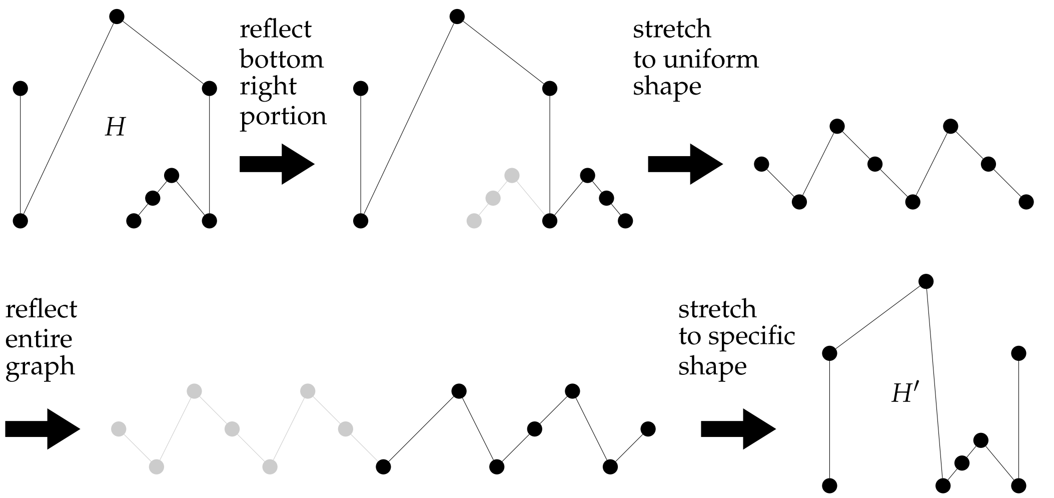

Figure 4. Demonstrating isomorphism between the causal graphs H and H' from Figure 3 via manipulations. Gray nodes and edges indicate "previous positions"; arrows indicate manipulative steps. Viewed as physical models, the two graphs are identical.

The automorphisms of a causal graph are its symmetries. In an obvious sense, these symmetries are analogous to the spacetime symmetries in special relativity defined by the Poincaré group, since causal graphs are themselves models of spacetime. However, in another sense, causal graph symmetries can share some similarities with the internal symmetries that characterize the standard model of particle theory, since many causal graph symmetries involve only the local permutations of a few nodes and edges. An important result obtained by Bender and Robinson [30] states that typical finite directed graphs are rigid, i.e., they have no nontrivial symmetries at all. Intuitively, this is because a typical graph is large, and large graphs offer more opportunities for symmetry breaking "defects", essentially for entropic reasons. In particular, any individual edge added to a symmetric graph can break the symmetry. However, this generic rigidity is not necessarily definitive for physical models. Assuming that discrete causal theory is valid at some level, the graphs that dominate the picture in representing physical spacetime must be very unlike typical graphs in a combinatorial sense. In particular, they must possess strong local properties, while typical graphs manifest the kind of small-world interconnectivity observed, for example, in social networks. Nevertheless, the intuition behind Bender and Robinson's rigidity result can be expected to apply in at least an approximate sense to most large-scale causal graph symmetries, i.e., those permuting large proportions of the nodes. Exact global symmetries seem to be very unlikely in most cases. For this reason, the apparent Poincaré symmetry of Minkowski spacetime, so intimately connected to the standard model, is likely to find its ultimate expression as an *emergent* approximate symmetry rather than an exact symmetry extending all the way down to the fundamental scale. Of course, Poincaré symmetry is not exact in *general* relativity either, since general relativistic spacetime is only approximately flat in local regions away from singularities. However, just as approximate local Poincaré symmetry is of central importance to small-scale behavior in continuum-based physics, the local symmetry properties of causal graphs are likely to prove crucial to understanding discrete causal physics near the fundamental scale.

The isomorphic causal graphs H and H' illustrated in Figures 3 and 4 are rigid; they possess no nontrivial automorphisms. The significance of this fact in the counterintuitive scenarios examined below is that each node in H or H' is qualitatively and uniquely *distinguishable* from all the other nodes by virtue of its structural relationship to other parts of the graph. For example, the top left node in H, labeled v_2, may be distinguished as "the unique second-generation node with empty future", where "second-generation" means that the longest directed path terminating at v_2 has a length of one. The top right node in H', labeled v_2' in Figure 3, possesses the same qualitative property, which is why it *must* be the image of v_2 under ϕ. The importance of such distinguishability lies in the fact that the identity-shifting nature of the scenarios we examine below cannot be explained away via symmetry. In fact, in Section 4.2, we demonstrate that a deeper *potential* symmetry is at work in this context.

2.4. Causal Graph Dynamics

The counterintuitive physical scenarios we wish to analyze involve causal graph dynamics, which describes how a classical history develops in discrete causal gravity. In rough terms, causal graph dynamics centers around the question of how a causal graph can "grow" by adding new nodes and edges to the future of existing nodes and edges. In physical terms, this translates to the details of how new events can be caused by previous events, and more generally how existing systems can evolve into future configurations. Since causal graph dynamics involves questions of "before" and "after", it is described in terms of natural relationships between pairs of causal graphs, i.e., morphisms $\phi : G \rightarrow H$, where the source graph G represents an early stage of spacetime evolution, while the target graph H represents a later stage. Not every morphism is permissible in this context; for example, it would make no physical sense to add nodes and edges "before" the existing structure of G, which is interpreted to represent the entire history of the system up to a certain point. We therefore begin by making precise the idea of a subgraph G of a causal graph H whose complement does not intersect with its past, and which can therefore serve as the image of a permissible morphism.

Definition 6. *A subgraph $G = (U, D, \alpha, \omega)$ of a directed graph $H = (V, E, \alpha, \omega)$ is called* **originary** *if its complement does not intersect its past; i.e., if there is no node $v \in V - U$ such that there exists a directed path γ from v to some node $u \in U$.*

For example, the full subgraph of H with node set $\{v_3, v_4, v_5, v_6\}$ in Figure 2 is an originary subgraph because none of the other nodes in H lie in its past. However, the full subgraph with node set $\{v_1, v_2, v_7, v_8\}$ is not originary because the node v_6 lies in its past. In physical terms, an originary subgraph G of a directed graph H may be viewed as an "earlier stage of evolution" of H; the nodes belonging to its complement represent events that are either "developing from" G, or "independent of" G.

A transition is a special kind of morphism between two directed graphs G and H that, naïvely, represents a legitimate evolutionary relationship between them. We denote transitions using the letter τ to distinguish them from general morphisms. In mathematical terms, a transition τ embeds G as a full originary subgraph of H. In this context, we generally identify G with its isomorphic image $\tau(G)$ in H. This means that H may be regarded as a "later stage of evolution" of G.

Definition 7. *A* **transition** *$\tau : G \to H$ between two causal graphs G and H is a monomorphism whose image $\tau(G)$ is a full originary subgraph of H. The graph G is called the* **source** *of the transition, and the graph H is called the* **target** *of the transition.*

Figure 5 illustrates transitions τ and τ' between a causal graph G and the isomorphic causal graphs H and H' shown together in Figure 3. The "new" nodes x and x' for each transition may be identified with the nodes v_8 and v'_5 in our original diagrams of H and H', and the "new" edges are illustrated by dashed lines. Also shown are a particular node y in G and its image nodes v_5 and v'_8 in H and H', which we later show to be pseudosimilar to the nodes $x = v_8$ and $x' = v'_5$, respectively. We return repeatedly to the properties of these transitions in what follows.

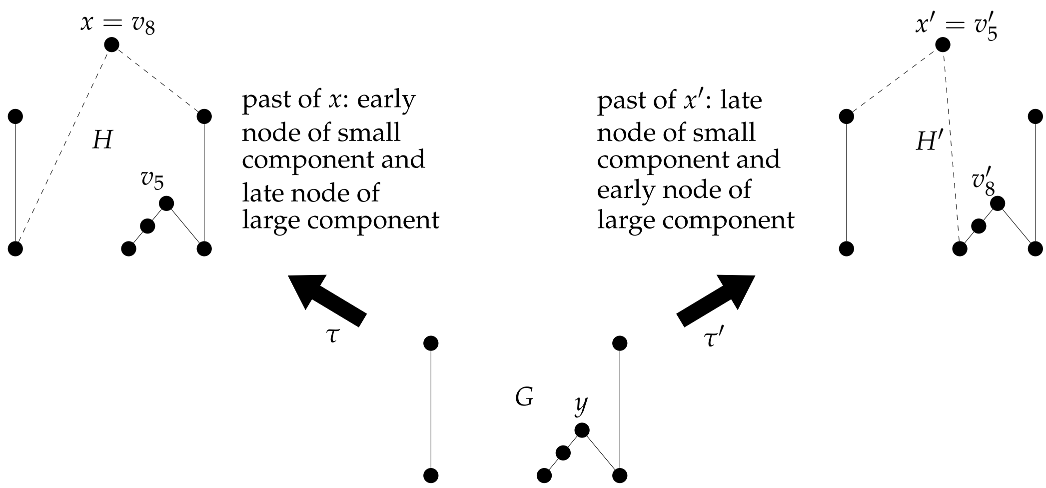

Figure 5. Transitions from a source graph G to isomorphic target graphs H and H'. The new node x in H does *not* map to the new node x' in H' under the isomorphism $\phi : H \to H'$. The pasts of x and x' are physically distinguishable. The images v_5 and v'_8 of the previously existing node y are pseudosimilar to x and x', respectively. "New" edges are indicated by dashed lines.

Just like any other morphism, a transition consists of a set map of nodes and a set map of edges. In some of the literature, transitions are, by definition, restricted to "adding only one event", in the sense that the size of the node set of the target H exceeds the size of the node set of the source G by exactly one. This is the viewpoint, for example, in classical sequential growth dynamics for causal sets [3], in which the term "sequential" indicates that nodes are being added one at a time. However, a broader definition for transitions can be useful when referring to evolutionary processes more generally, since each "time step" might be naïvely expected to add an entire "slice" or "generation" of nodes to the developing system.

3. Relativity of Identity

3.1. Pseudosimilarity

The reasons why transitions only *naïvely* represent evolutionary relationships between pairs of causal graphs arise from a pair of interrelated complications. The first is that physical significance is ascribed only to *isomorphism classes* of causal graphs in discrete causal gravity, rather than to specific representatives of these isomorphism classes. This is because it is *causal structure itself* that we consider to be physically significant, not the details of how nodes and edges are labeled or drawn. Labeling nodes and edges is often necessary for analytical purposes, and is analogous to a choice of gauge in familiar physical theories such as electrodynamics. However, no physically meaningful quantity can depend in any essential way on such labels. The reason this introduces a complication in the context of transitions is because transitions are defined a priori as morphisms between specific pairs of causal graphs, and care is required to translate this definition into a physically meaningful one in terms of isomorphism classes. The second complication is closely connected to the first, but much subtler. It turns out that *a given pair of causal graphs can be related by multiple transitions in essentially different ways*, and this raises deep questions about their physical interpretation. This phenomenon requires careful analysis, first to establish that the counterintuitive behavior involved cannot be explained away via conceptual devices already familiar in modern physics, and second to analyze what is actually happening and how it should be interpreted. In particular, we establish that such scenarios *cannot* be explained away via symmetry, the relativity of simultaneity, gauge-like relationships, or any other familiar notion. They involve a graph-theoretic property called pseudosimilarity that is absolutely foreign to continuum-based models of spacetime. We begin by defining and describing this property, then go on to analyze its physical significance.

The graph-theoretic literature contains a broad variety of different notions of pseudosimilarity, most of which are far removed from our present setting. We therefore adapt definitions appropriate for our purposes without working from any specific source. The simplest version of pseudosimilarity involves pairs of individual nodes in a directed graph H. Such nodes can be similar to each other in one way and dissimilar in another way, and the tension between these two types of similarity leads to surprising physical consequences. The way in which such nodes are similar is that removing them, together with all edges connected to them, from two copies of H, leaves isomorphic subgraphs. From a physical perspective, this means that the hypothetical histories constructed by deleting these events are identical. One would expect, then, that the two events themselves would play identical roles in the history represented by H, but they do *not*: there is no symmetry of H, i.e., no automorphism, that interchanges the two nodes.

Definition 8. *Nodes x and y in a directed graph H are called* **pseudosimilar** *if the complements of x and y in H are isomorphic but there exists no automorphism of H interchanging x and y.*

As demonstrated in Section 4.1 below, the notion of pseudosimilarity can be extended to pairs of subgraphs of H, which may be arbitrarily complex. For the present, however, we explore the consequences of pseudosimilarity for pairs of individual nodes.

3.2. Passive Viewpoint and Failure of Cancellation

Pseudosimilarity interacts with transitions in a subtle way that involves the discrete causal analogue of the relativity of simultaneity, but leads to consequences far beyond those that arise in ordinary relativity. An important element in analyzing this relationship is the contrast between "active" and "passive" viewpoints regarding a transition. For example, in the case where the two graphs differ by a single node, one may choose to think of actively adding a node x, together with some edges, to a graph G to obtain a larger graph H, or one may choose to think of passively identifying G as a subgraph of H via a transition. We examine the passive viewpoint first because it explicitly involves transitions and is easier to relate to the definition of pseudosimilarity. Figure 6 shows two different transitions between the same pair of causal graphs G and H, involving pseudosimilar nodes identified as the nodes v_5 and v_8 in the labeled diagram of H in Figure 2. The source graph G consists of two connected components with two and five nodes, respectively, colored red and blue for ease of identification. The red and blue subgraphs of the target graph H indicate the images of these two components under the transitions τ and τ'. The transition τ appears "obvious" according to how the graphs are drawn, because the image $\tau(G)$ looks identical to the source G. The transition τ' appears "non-obvious" in the sense that its image $\tau'(G)$ does not look like G at first glance, but the two are easily verified to be isomorphic. The temptation to regard τ as "more valid" than τ' is similar to the temptation to regard frames of reference represented by orthogonal axes in Minkowski spacetime diagrams as "more valid" than frames represented by "scissored-together" axes. The alternative between the equally valid viewpoints represented by these different-looking subgraphs is analogous to the familiar relativity of simultaneity, and is *not* the issue at stake. The issue, rather, is that the two transitions represent physically distinguishable, irreducible alternatives for how to add the "new" node, yet result in *exactly the same history H.*

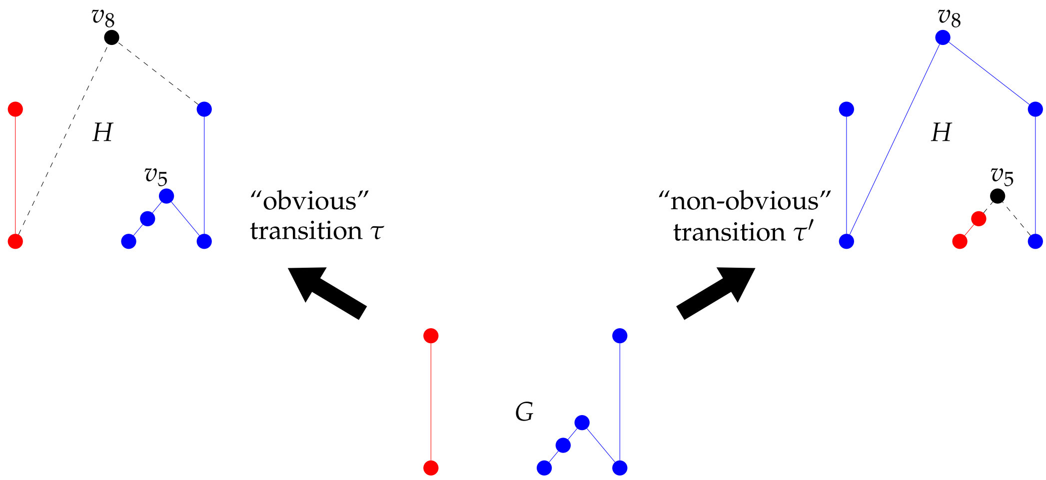

Figure 6. Pair of transitions involving pseudosimilar nodes v_8 and v_5: either node may be considered "new" by the relativity of simultaneity, but no automorphism of H interchanges them. Colors distinguish the images of the smaller (red) and larger (blue) components of G in H under τ and τ'.

For context, we now consider two simpler scenarios that lack the full consequences of pseudosimilarity. Figure 7A shows two possible transitions between two simpler graphs labeled G and H, but these transitions do not really represent different physical scenarios because the target graph H has a symmetry interchanging the two second-generation nodes. The *same* physical description that "the earlier event a in the history G produces a new event unrelated to the other event b produced by a" therefore applies to both transitions. By contrast, the two transitions illustrated in Figure 7B between G and a different target graph H' exhibit some of the features of pseudosimilarity: the pairs of "new" nodes added in each case, taken together, are physically distinguishable, lead to isomorphic subgraphs of H' when they are deleted, and are not interchanged by any automorphism of H'. However, the process is reducible in this case because the two nodes may be added one at a time, and the counterintuitive features disappear when the process is broken down in this way, regardless of which choices are made about order. Indeed, choosing to first add a node emanating from the node a reproduces the same scenario as in Figure 7A, followed by the physically indistinguishable choice between two second-generation nodes as the source of the final node. Alternatively, choosing to first add a node emanating from the node b leaves no ambiguity about adding the final node.

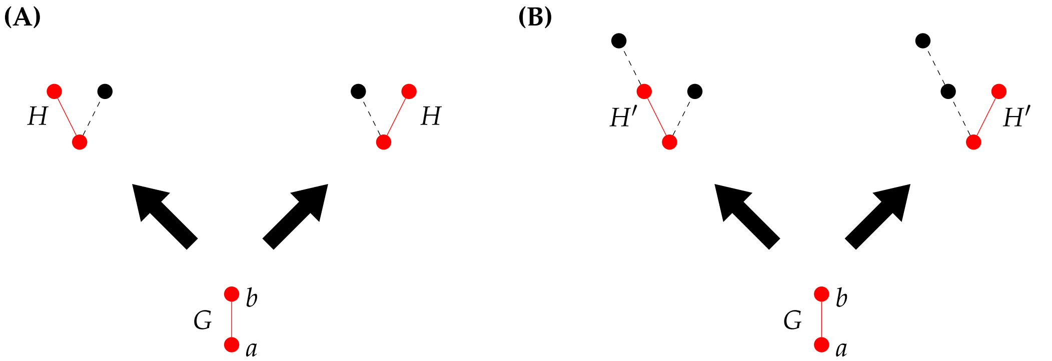

Figure 7. (**A**) Transition whose counterintuitive features can be explained away by symmetry; (**B**) transition whose counterintuitive features are ameliorated by reducibility.

The scenarios shown in Figures 6 and 7B both violate a very basic expectation about the consequences of distinguishable alternative events. Intuitively, it is natural to expect that, if one begins with two identical systems with identical histories, then observes physically distinguishable alternatives transpiring within each system, one will obtain distinguishable future histories. The violation of this basic expectation due to pseudosimilarity may be summarized by the following heuristic equations:

$$G + (\text{new events } A) = G + (\text{new events } B) \qquad \textit{even though} \\ A \neq B \tag{1}$$

This phenomenon may be viewed as a breakdown of a sort of "cancellative property" for histories: one would expect that G could be canceled from both sides, leaving $A = B$. The failure of this expectation in the reducible scenario from Figure 7B is less surprising, because when multiple events are involved, it is within the bounds of ordinary intuition to expect that the effects of different events could balance each other out in such a way as to allow for the same overall outcome. However, the strangeness of the scenario in Figure 6 is unavoidable.

3.3. Active Viewpoint and Relativity of Identity

The active viewpoint for transitions involving pseudosimilar nodes presents, if possible, an even stranger aspect than the passive viewpoint. From the active viewpoint, one or more new nodes are added to copies of an existing history G in different ways. To understand this viewpoint, we return to Figure 5. In this case, both transitions are equally "obvious" in the sense that their images appear identical to the source graph G. What is *not* obvious is that the two target histories H and H' are isomorphic, but we have already demonstrated this isomorphism in Figures 3 and 4. The alternatives represented by the two transitions are clearly physically distinguishable: for τ, the new node x emanates from the initial node of the smaller connected component and a second-generation node of the larger component of G, while for τ', the new node x' emanates from the terminal node of the smaller connected component and an initial node of the larger component. These descriptions are qualitatively different, yet the resulting histories are *physically identical* since the graphs are isomorphic! How is this possible? Referring back to Figure 3, the subtlety is that the "new" node, identified with v_8 in H and v'_5 in H', *swaps physical roles* with the "previously existing" node, identified with v_5 in H and v'_8 in H', depending on which transition occurs. This interchange of roles, which we call relativity of identity, is formalized by the fact that the "old" and "'new" nodes are interchanged under the isomorphism ϕ illustrated in Figure 3. Relativity of identity "balances" or "cancels out" the physically distinguishable alternatives about how the new node is added, thereby leading to identical histories. Retrospectively, one may apply the familiar notion of relativity of simultaneity to these nodes: either node may be regarded as having occurred first. But the novel issue is *which node is which* physically. The answer, of course, is that the node labeled v_8 in H must be retrospectively identified with the node labeled v_8' in H', due to the equivalent causal roles played by these nodes in the target histories. However, in the source history G, the node y in Figure 5 mapping to v_5 under τ and v_8' under τ' is uniquely identifiable because G is rigid. Nor do the images of y lose identifiability in the isomorphic target histories H and H', which are also rigid. It seems unavoidable to conclude that uniquely identifiable components of identical systems with identical histories can acquire *different* future identities in some sense, even if the overall futures are identical.

Before proceeding, we dispose of an elementary source of confusion involving the distinction between identical configurations and identical histories. It is not strange or novel, at least from a logical perspective, for two systems that are different at one point in time to eventually become identical (convergence), or for two initially identical systems to eventually become different (divergence), although such outcomes may clash with conservation or reversibility or uniqueness principles in certain special cases. However, the scenarios we are presently considering involve the *entire histories of systems that are causally closed.* In particular, the systems represented by the graphs H and H' in Figure 5 are not only identical, but always were identical in the most absolute sense. This is true despite the fact that, at the earlier stages of evolution of these systems represented by the transitions τ and τ', physically different alternatives occurred, and therein lies the oddity. Considering again the relativity of simultaneity, one could of course select two *nonisomorphic* earlier stages of evolution of a typical history H, by simply deleting different parts of its late-term structure. The resulting subgraphs could then be built back into H by adding the appropriate structure to each. More generally, given any pair of causal graphs, one may build a common future in which both are originary subgraphs, thereby defining a sort of "convergence of histories". The present case, however, is different: identical systems with identical histories, subject to different events, produce identical results, which were always identical in every respect.

3.4. Analogies Involving Macroscopic Events

We now consider how such scenarios might appear if they were to occur at the macroscopic level involving ordinary objects and events. Thorough understanding requires the consideration of both the active viewpoint, in which subsequent events are added to an existing history G to obtain a larger history H or H', and the passive viewpoint, in which G is recognized as the image of a transition into H or H'. The active viewpoint conforms most closely to our basic intuition about the existence and non-alterability of G, while the passive viewpoint offers a clear a posteriori view of how H or H' could have developed from G. Figure 8 depicts two alternative idealized scenarios involving a cat and a dog and the effect of their behaviors on a human observer. In Figure 8A, the cat takes some action early in its existence that has an effect on the human, while the dog waits until later before causing similar trouble. In Figure 8B, the roles are reversed: the early behavior of the dog and the late behavior of the cat do the damage. Few would seriously consider the possibility that these two scenarios might represent the *same history*; the distinctions between them are self-evident. The key feature is that the cat and dog are distinguishable. By contrast, if each scenario were to involve two cats, it would be at least logically possible that they could be identical, and that spacetime and all its contents as a whole could possess a symmetry under which they could be interchanged. In this unlikely but not impossible case, the two scenarios might be identified as different points of view regarding the same history. However, the distinguishability of the cat and the dog absolutely rules out any such consideration.

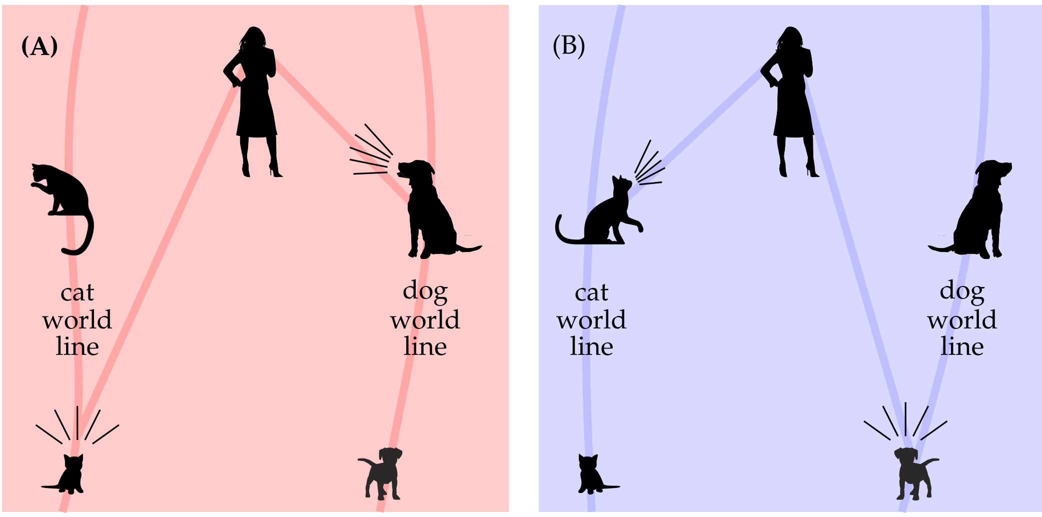

Figure 8. Alternative histories: (**A**) kitten mews, then dog barks; (**B**) puppy whines, then cat meows.

Or does it? Figure 9 shows an abstract version of the same scenario, which is readily recognized as the same pair of inequivalent transitions between a source graph G and the isomorphic target graphs H and H' previously shown in Figure 5. As in the case of the macroscopic cat and the dog, the histories of the new events in H and H' are physically distinguishable. The fact that they are small and abstract is beside the point; they possess the same property of qualitatively distinct features as the macroscopic cat and dog. Using the macroscopic analogy as a guide, we now re-examine the implications of the isomorphism $\phi : H \to H'$. In H, the node v_1 represents the "kitten", while in H', the corresponding node $v'_1 = \phi(v_1)$ represents part of the "puppy". Similarly, in H, the node v_2 represents the "cat", while in H', the corresponding node $v'_2 = \phi(v_2)$ represents the "dog". The isomorphism ϕ therefore mixes the identities of the "kitten" and "puppy", and completely swaps the identities of the "cat" and the "dog". This occurs despite the fact that we can unambiguously identify the "kitten/cat" as the smaller connected component and the "puppy/dog" as the larger connected component in *both* histories before either of the "new" nodes are added. It seems, in some sense, that the "human observer" *determines which "animal" is which* via her observations.

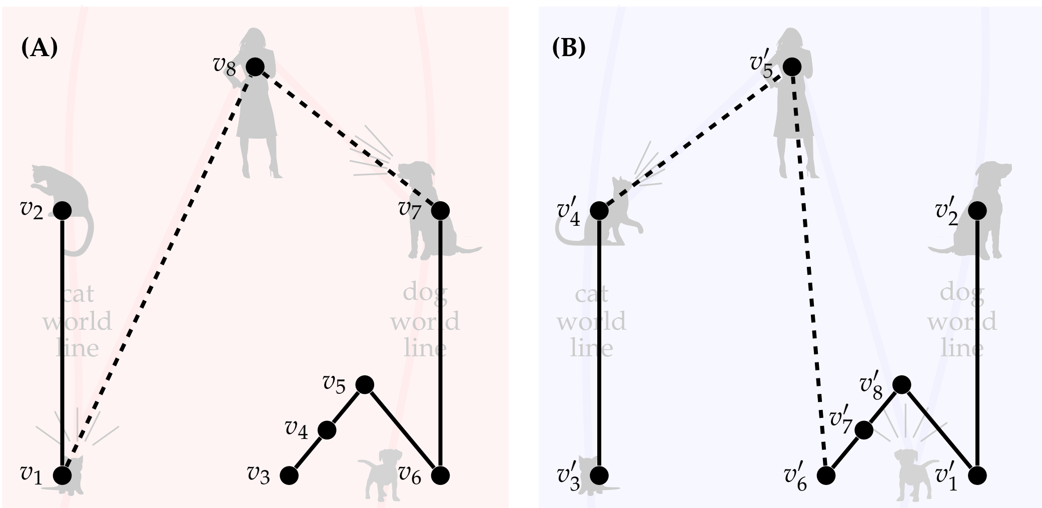

Figure 9. Discrete causal analogue of cat/dog histories: (**A**) new event caused by the initial event of a small component and the later event of a large component; (**B**) new event caused by the initial event of a large component and the later event of a small component. Event origins are physically distinguishable but the histories are isomorphic!

The strangeness extends further. In H, the node v_8 represents the "human observer", while in H', the corresponding node $v'_8 = \phi(v_8)$ represents part of the "puppy". Meanwhile, the node v_5, pseudosimilar to the node v_8 in H, represents part of the "puppy", while the node v'_5, pseudosimilar to the node v'_8 in H', represents the "human observer". As in a Bach fugue, where the voices of the instruments assume different identities in different contexts, the structural components of G, though initially distinguishable in an unambiguous way, swap and/or mix identities with each other and/or with their complements in H and H'.

3.5. Relativity of Identity as an Extension of the Relativity Principle

If there is any existing precedent for such a behavior in modern physics, it involves mechanisms such as entanglement, which are intrinsically quantum-theoretic. However, the present context is purely classical. What we seem to be encountering is a radical and unanticipated extension of the relativity principle beyond the now-familiar relativity of spatial and temporal measurements in Einstein's theory to encompass the relativity of the identities of different components of a causal structure, including discrete causal classical spacetime as our featured special case. To place this phenomenon in its proper context, we explore how it compares to other notions of relativity, then proceed to further analyze how it applies to spacetime, physical processes, and observers.

In referring to the "relativity principle" in a general qualitative sense, we mean the overarching paradigm shift away from absolute, uniform notions of physical quantities and/or qualities, towards observer-dependent or contextual notions, principally driven by the theories of special and general relativity. The relativity principle manifests itself via a broad variety of more-specific statements about particular quantities or processes. For context, we review a few of its more familiar manifestations, most of which appear in some form in special relativity, with further modifications due to gravitational effects on spacetime structure in general relativity:

1. **Length contraction:** The observed length of a physical object in special relativity is not absolute, but dependent upon the state of motion between the object and the observer. General relativity adds further modifications.
2. **Time dilation:** The observed time interval between two events in special relativity is not absolute, but dependent upon the world line traced out by the observer between the two events. General relativity adds further modifications.
3. **Relativity of simultaneity:** The observed order in which spacelike-separated events occur in special relativity is not absolute, but dependent upon the observer's frame of reference. General relativity adds further modifications.
4. **Relativity of dynamical quantities:** The observed mass, energy, momentum, and other such quantities associated with a physical object in special relativity are not absolute, but dependent upon the state of motion between the object and the observer. General relativity adds further modifications.
5. **Relativity of spacetime structure:** Spacetime structure in general relativity is not uniform, but dependent upon the distribution of matter and energy determining the gravitational field, including the matter and energy associated with observers and observational equipment.

Into this tableau fits the relativity of identity as a postmodern "digital" addendum, arising from the information-theoretic properties of discrete causal spacetime models:

6. **Relativity of identity:** The observed identities of the components of a physical system in discrete causal theory, including spacetime itself, are not absolute, but dependent upon the causal relationship between the system and the observer.

We emphasize once more the radical degree of disagreement that can arise between different observers under the relativity of identity, which qualitatively exceeds any conceivable disagreement about length or time measurements. This disagreement reflects equally valid but drastically different points of view. It cannot be resolved by symmetry because the components of the systems involved are absolutely distinguishable. It cannot be resolved by gauge-like considerations because it does not depend on labels or other nonphysical auxiliary quantities. It cannot be resolved by examining a deeper level of detail because it occurs at the fundamental scale. It cannot be resolved by convergence or divergence, because it involves the entire history of the observed system. As an extreme case, it implies the theoretical possibility that different astronomers could legitimately disagree about *which parts of the universe* certain uniquely identifiable signals originate from. For example, if the nodes v_5 and v_8 in the graph H in Figure 6 represent "astronomers", then each astronomer believes *herself* to be the one observing the smaller (red) component of the early universe G, in a "no, *I* am Spartacus" fashion. As a final illustration, we offer the macroscopic analogy of Green Science Hall: under previous notions of relativity, different observers could legitimately disagree about the length of the building, the time elapsed between a meeting in Room 108 and a meeting in Room 204, who entered the building first, the maximum load carried by the elevator, the deviation of the main hallway from true north, and so on. However, under the relativity of identity, observers could disagree about which department is the Math Department and which is the Biology Department!

4. Generalizations and Discussion

4.1. Generalized Relativity of Identity

Before further analyzing the implications of the relativity of identity, we first demonstrate the ubiquity of such scenarios and analyze some of their symmetry-based origins. Our first task is to show that the specific scenario involving the directed graphs G, H, and H' analyzed above is not a unique oddity. In fact, the identity of *almost any* finite directed graph, embedded as a subgraph of an appropriate history, may be interchanged with the identity of a disjoint subgraph of this history via the addition of alternative versions of appropriate future structure to form larger histories. Technically, this interchange means that any isomorphism of the larger histories will map the original graph to a disjoint subgraph. We now demonstrate this striking fact via the following Theorem 1, illustrated in Figure 10. To streamline the proof, we use the terminology of attaching a directed graph G'_1 "to the future" of a directed graph G to mean that each attaching edge begins at a node in G and terminates at a node in G'_1. We use the terminology that G'_1 may be attached to G in two or more "nonisomorphic ways" if no automorphism of G interchanges the families of nodes involved in the two attachments. We extend this terminology in the obvious way to apply to the case of attaching G'_1 to two different copies of G: the attachments are considered nonisomorphic if there is no isomorphism between the two copies mapping the family of nodes involved in one attachment to the family of nodes involved in the other attachment.

Theorem 1. *Let G_1 be a finite directed graph with at least one pair of nodes that are not interchanged by any automorphism of G_1. Then there exists a finite directed graph G containing G_1 as a subgraph and satisfying the following properties:*

1. *A directed graph G'_1 may be attached to the future of G in two nonisomorphic ways such that there is an isomorphism ϕ between the resulting directed graphs H and H', where the image $\phi(G_1)$ is a subgraph of G, disjoint from G_1, when both are viewed as subgraphs of H'.*
2. *No isomorphism between H and H' fixes G_1, viewed as a subgraph of G.*

Proof. Referring to the schematic drawing in Figure 10A, let G_1 be a finite directed graph satisfying the given hypothesis, and let G'_1 be any other finite directed graph. Construct a graph G by first cloning G_1 twice to obtain a total of three disjoint copies G_1, G_2, G_3, then cloning G'_1 and attaching the resulting graph G'_2 to the futures of the graphs G_2 and G_3 in nonisomorphic ways. This is possible due to the hypothesis on G_1. Solid lines in the figure represent the attachments of G'_2 to G_2 and G_3, which may in fact consist of many edges each. Square and triangular "attachment points" on the schematic drawings of G_1, G_2, G_3 indicate that these attachments are made in nonisomorphic ways. The other ends of these attachments, at G'_2, may be the same, as suggested by the figure, or may be different, without affecting the result. G, not explicitly labeled, is the resulting graph consisting of the four subgraphs G_1, G_2, G_3, and G_2', with the latter three graphs attached together. To construct the graph H, attach G_1' to the future of G_1 in the same way that G_2' is attached to the future of G_2, and attach G_1' to the future of G_2 in the same way that G_2' is attached to the future of G_3, as shown in the figure.

To construct the graph H', shown schematically in Figure 10B, first repeat the previous construction of G. Then, attach G_1' to the future of G_1 in the same way that G_2' is attached to the future of G_3, and attach G_1' to the future of G_3 in the same way that G_2' is attached to the future of G_2. The graphs H and H' then admit an isomorphism mapping the copy of G_1 in H (pink) to the copy of G_2 in H' (pink), which are disjoint in the common subgraph G of H and H'. This isomorphism also maps the copy of G_2 in H (green) to the copy of G_3 in H' (green), and the copy of G_3 in H (blue) to the copy of G_1 in H' (blue), while swapping the copies of G_1' and G_2'. No isomorphism between H and H' fixes G_1 because the attachments between G_1' and G_1 in H and H' are nonisomorphic. □

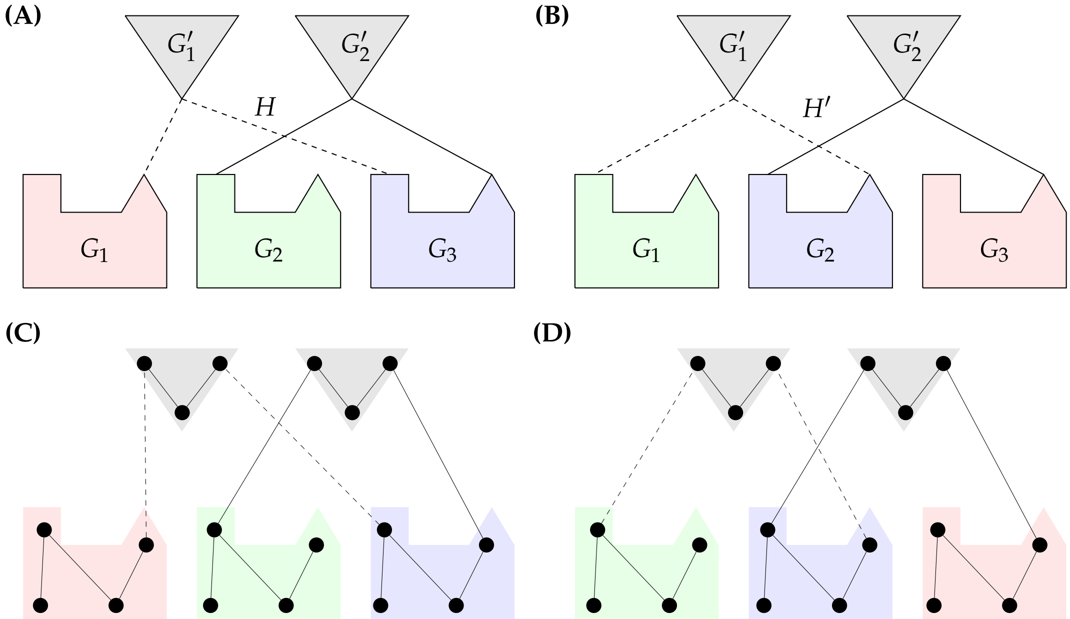

Figure 10. Generalized relativity of identity: (**A**,**B**) the graph G_1' may be added in two distinct ways to the graph constructed from G_1, G_2, G_3, and G_2'; the resulting graphs H and H' are isomorphic, but the identities of the subgraphs G_i, G_j' are permuted; (**C**,**D**) concrete example of this generalized construction.

From an abstract perspective, the entire subgraphs G_1' and G_2' are pseudosimilar in H, and likewise in H', in the sense that deleting them produces isomorphic subgraphs, but no automorphism of H or H' can interchange them, since they attach in nonisomorphic ways to a common copy of G_1. However, as mentioned in Section 3.2, the case in which G_1' and G_2' consist of individual nodes is in some ways the most interesting, because it renders the process irreducible and thereby prevents the strange properties involved from being explained away via a series of intermediate steps. Figure 10C,D illustrate an example of this construction, with the schematic shapes included in the background as a visual aid. Other such constructions are possible; for example, one may begin with *four* copies of a graph G_1 admitting at least three nonisomorphic attachments, then attach copies of another graph G_1' to three copies of G_1 each. From a physical perspective, this generalization demonstrates that *arbitrarily complex objects*, e.g., as complex as cats and dogs, can assume different identities under alternative additions of future structure in the discrete causal context. Furthermore, there is nothing about these constructions from a mathematical perspective that requires the copies of G_1' to be attached to the *future* of the copies of G_1; for example, one could immediately generate "dual constructions" by simply "flipping the page". Similarly, these constructions could be immediately generalized to admit the existence of causal cycles without changing their basic implications regarding the relativity of identity.

4.2. Symmetry Explanation

As in so many other areas of modern physics, symmetry considerations may be found at the heart of the pseudosimilarity-related constructions in our present study. The resulting physical notion of relativity of identity is therefore intimately connected to symmetry properties, just as more-familiar manifestations of the relativity principle are connected to the symmetries of the Poincaré group. The present connection is perhaps less obvious, since the definition of pseudosimilarity hinges on the *absence* of a symmetry interchanging the nodes or subgraphs in question. Upon reflection, however, we recognize a deeper underlying notion of symmetry at work, involving an additional *potential* structure whose properties may be illustrated in a particularly pleasing way. We demonstrate this symmetry-based explanation of the relativity of identity for the particular graphs G, H, and H' first introduced together in Figure 5, which have played such an important part throughout our study. In this case, the graph G_1 in Theorem 1 consists of just two nodes, vertically connected by an edge. Multiple copies of this graph, color-coded for ease of reference, appear in the schematic cylinders shown in Figure 11A,B. Meanwhile, the graph G'_1 consists of just a single node in this case; multiple copies of this graph appear at the tops of the cylinders. Referring back to Figure 10A, the nonisomorphic "attachment points" on G_1 represented by the square and triangular shapes are in the present case just the bottom and top nodes of G_1, respectively.

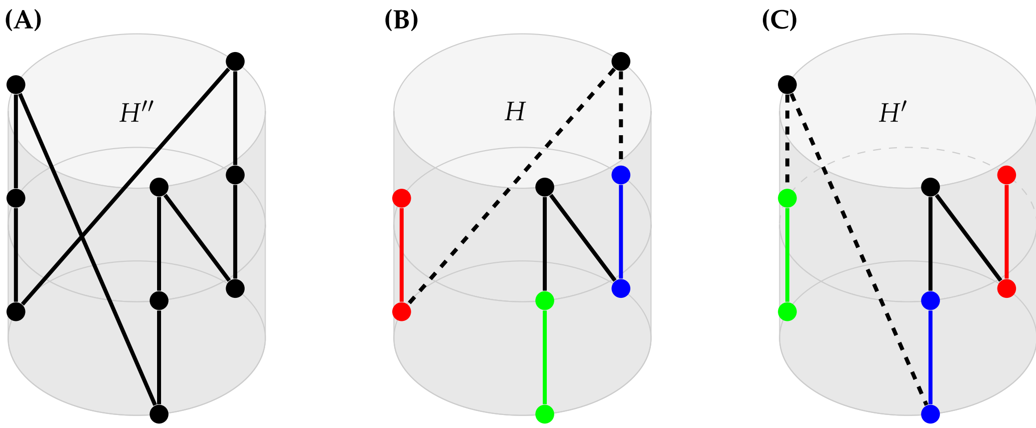

Figure 11. (**A**) history H'' with $\mathbb{Z}_3$ symmetry; (**B**,**C**) isomorphic sub-histories H and H' without symmetry, featuring pseudosimilar events, which are generated in physically distinguishable ways. Dashed lines indicate "new" events, in reference to Figure 5. Colors indicate how this construction is a special case of the general construction illustrated in Figure 10.

Extending the construction in the proof of Theorem 1, we now introduce a *third* copy G'_3 of the graph G'_1, which we attach to the futures of G_1 and G_2 at the remaining open attachment points. The resulting graph H'' now possesses a $\mathbb{Z}_3$-symmetry group corresponding to the cyclic permutations of the indices of the subgraphs G_i and G'_j, and the isomorphism $\phi : H \to H'$ is a transparent consequence of the corresponding group action. In the present case, for example, ϕ is represented by rotating the cylinder clockwise through an angle of $\pi/3$. The clarity of this viewpoint, contrasted with the "non-obvious" appearance of ϕ in Figure 3, suggests that the notion of potential symmetry offers a valuable perspective on pseudosimilarity and the relativity of identity. From a physical perspective, the additional structure of G'_3 required to produce the symmetric history H'' is entirely hypothetical. For example, returning to the macroscopic analogy of the cat and dog, this structure might correspond to a second human observer affected by the opposite alternative pair of events. This seems to imply that the stability of the identities of physical objects could depend upon *potential* symmetries that could hypothetically be generated by future events, another counterintuitive physical consequence of pseudosimilarity.

4.3. Discussion and Conclusions

Before discussing some further topics involving discrete causal quantum theory and phenomenology, we briefly revisit the implications of our study for several established approaches to quantum gravity. For simplicity and definiteness, we focus on two theories incorporating discrete causal assumptions at a fundamental level, namely causal set theory and causal dynamical triangulations, and two theories for which notions of spacetime discreteness are derivative, namely loop quantum gravity and superstring theory. We first emphasize that *any* theory involving causal structure, even Newtonian mechanics, includes the possibility of functional or approximate relativity of identity, in the sense that such considerations will always arise whenever we focus exclusively on causal structure and ignore other means of distinguishing the components of physical systems. However, our principal focus throughout our present study has been the *absolute* relativity of identity at the fundamental spacetime level.

For causal set theory, the application of these ideas is immediate and straightforward, because the directed graphs H and H' introduced in Figures 3 and 4 may be *converted* into causal sets by closing them under transitivity, i.e., by adding edges connecting the node pairs (v_3, v_5) and (v_6, v_8) in H, and adding corresponding edges in H'. This modification preserves all the necessary relationships between the graphs and their constituent parts; in the language of [10], transitive closure is a functor. In particular, the same nodes remain pseudosimilar in the two graphs, and the graphs themselves remain isomorphic. Similar considerations apply to most of the other causal graphs considered in previous sections. Our analysis of the relativity of identity therefore carries over to causal set theory with little modification.

Pseudosimilarity also arises in causal dynamical triangulations, although with some practical restrictions on its effects, as we now describe. In causal dynamical triangulations, the nodes representing discrete spacetime units in our general causal graph approach are replaced by "4-simplices" glued together in specific ways. In the standard case, these simplices carry intrinsic four-dimensional Lorentzian structure, although different effective dimensions can arise dynamically. More generally, one may work with n-simplices, since other spacetime dimensions are sometimes worth considering. It is straightforward to construct triangulations whose underlying graphs, obtained by collapsing each simplex to a node, are isomorphic to graphs with properties like those of H and H'. However, a symmetry-breaking complication arises from the fact that different simplices are often intrinsically distinguishable by counting how many of their vertices reside at different causal levels in the triangulation. Specific rules for each dimension govern which types of simplices may be glued together, since the corresponding faces must match. This restricts the family of triangulations that exhibit "absolute pseudosimilarity" in an appropriate sense, i.e., when taking all the structure into account. Relativity of identity still occurs under certain conditions, but precisely characterizing its scope would require significant further analysis.

The role of pseudosimilarity in loop quantum gravity and superstring theory is less certain since discrete causal structure is less fundamental to these theories. Loop quantum gravity involves discrete spatial structures called spin networks, which are elevated to incorporate causal structure as spin foams. As in the case of causal dynamical triangulations, one may identify underlying graph-theoretic structure, with the potential for pseudosimilarity and consequent relativity of identity. However, significantly more auxiliary structure is present in this case, offering more opportunities for symmetry breaking. For example, the edges of spin networks are labeled by group representations, so different edges are generally inequivalent. The explicit inclusion of spatial structure itself does not preclude the relativity of identity; for example, the "spatial structure" provided by the schematic cylinders in Figure 11 merely elucidates the causal relationships involved. However, precisely characterizing the possible role of pseudosimilarity in this context would require significant further analysis. In superstring theory, discrete causal structure is much less prominent, and no clear graph-theoretic structure is present at the fundamental scale. Hints of discreteness, or at least some of its properties, arise from the minimal size imposed by the string scale. This is enhanced by T-duality, which implies the equivalence of reciprocal scales under certain conditions. It is not clear, however, that any straightforward mechanism exists to generate fundamental pseudosimilarity in this context.

Although the relativity of identity arises at the classical level in discrete causal gravity, its consequences unavoidably affect the sum-over-histories approach to the corresponding quantum theory. In this context, all possible evolutionary relationships among all possible classical histories must be considered simultaneously. Dynamics is determined by assigning phases to evolutionary pathways through the resulting structured configuration space, in a manner analogous to Feynman's phases for particle trajectories through a manifold. Figure 12 illustrates a small part of this configuration space structure, called the "positive sequential kinematic scheme" in [10], which is a modified version of Rideout and Sorkin's configuration space for sequential growth dynamics of causal sets [3]. Each large node in the figure contains a causal graph, considered up to isomorphism class. Causal graphs more general than causal sets are indicated by the gray nodes in the figure. The sequential development of any classical history may be described via an upward-directed path through this configuration space, relying on the useful fact that the entire space inherits a natural "up the page" direction for evolutionary relationships from the corresponding structure of its individual causal graphs. The large-font numbers indicate the numbers of nodes in each graph, increasing sequentially along any upward-directed path.

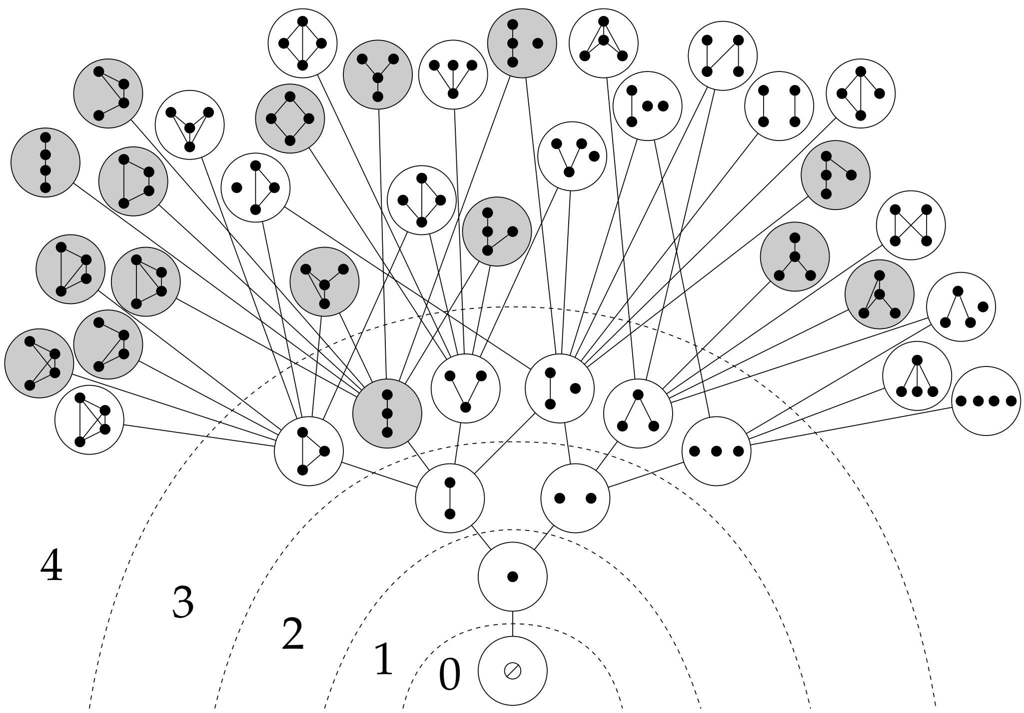

Figure 12. Quantum gravity via path summation over a configuration space of causal graphs. The empty history is denoted by ⊘. Large numbers indicate numbers of nodes. Gray shading indicates histories that are not causal sets.

Relativity of identity enters the picture via pseudosimilarity when we attempt to specify exactly what the edges connecting the large nodes should represent mathematically. Physically, they are intended to stand for "individual sequential evolutionary relationships", but it is not a priori obvious how to capture this idea in terms of causal graphs. Transitions seem to be the wrong answer, since multiple physically equivalent single-node transitions are possible between the same pair of causal graphs due to symmetry, as illustrated in Figure 7A. Should an edge then represent the existence of *at least one* transition between the two graphs? Again, the answer seems to be no, because multiple physically *inequivalent* single-node transitions are also possible between the same pair of causal graphs due to pseudosimilarity, as illustrated in Figure 5. In [10], this issue is addressed by grouping physically equivalent transitions together into "co-relative histories", leading to a configuration space with multiple edges between certain pairs of causal graphs. For example, the isomorphism class of the graph G in Figure 5 and the common isomorphism class of the graphs H and H' are related by two different co-relative histories. The importance of these details is due to the distinct roles of physically equivalent and inequivalent transitions in determining the correct choices for the phases assigned to such edges in the quantum theory.

Since observation informs intuition, the apparent strangeness of phenomena such as relativity of identity is enhanced by unfamiliarity. For example, we do not in fact observe the species of animals to be dependent upon our interactions with them. Hence, it is instructive to examine why relativity of identity does not seem to play an obvious role in macroscopic events. For context, we first review why some of the standard named "paradoxes" mentioned in Section 1 escape ordinary observation and therefore strike our intuition as paradoxical. For Class A scenarios, which involve well-understood physics, explanations are usually straightforward and center on the unfamiliarity of the scales, energies, relative velocities, or precisions at which these scenarios are expected or known to occur. For example, the twin paradox falls outside everyday experience simply because human civilization currently lacks the necessary mastery of propulsion to enable sufficiently accelerated travel. As explained in Section 1, however, essentially equivalent behavior is observed routinely in particle physics, and there is no serious doubt that the same basic phenomenon would manifest itself if the necessary accelerations could be applied to macroscopic objects. Similar considerations apply to the barn–ladder paradox: we currently lack the ability to directly compare significant Lorentz contraction from the rest frames of macroscopic objects in fast relative motion, since this would require the same type of relativistic propulsion technology necessary to test a full-scale version of the twin paradox. However, there is again no serious doubt about the nature of the phenomenon, due to the essentially equivalent behavior that we can observe.

Class B scenarios are generally murkier by definition, partly because it can be unclear what we should be trying to explain. In particular, such scenarios are sometimes regarded not as potentially real phenomena to seek out via experimentation, but as evidence that something is wrong with the theory. According to such a "pessimistic" viewpoint, the explanation of why these phenomena are not observed is simply that they do not in fact happen, and the focus of explanation shifts to analyzing where the theory might be breaking down and how it should be modified. If, on the other hand, we trust the theory sufficiently, then we are prompted to undertake the same type of "optimistic" analysis applied to Class A scenarios, namely, to provide a mechanistic explanation of how the phenomena can occur and yet "hide from view". For the example of Schrödinger's cat, pessimistic viewpoints posit that quantum theory is simply wrong and should be replaced with another theory such as a stochastic or hidden variable theory. Optimistic viewpoints assume quantum-theoretic superposition to be valid, prompting the further analysis of why we fail to observe *macroscopic* superpositions, which then leads one to the consideration of possible mechanisms such as collapse of the wave function, many worlds, decoherence, and so on. For the grandfather paradox, pessimistic viewpoints rule out closed causal curves entirely, for example, by invoking quantum effects. Optimistic viewpoints admit the possibility of such curves, prompting further analysis of how they could produce consistent consequences.

For the relativity of identity, any pessimistic viewpoint would be more or less obliged to reject or drastically restrict the discrete causal paradigm itself, since such scenarios will inevitably arise for any sufficiently general class of causal graphs. Such pessimistic arguments could be based on optical experiments; in particular, specific types of lattice-based discrete spacetime models have been criticized on such grounds. Causal set-type models involving an irregular or random structure seem to remain quite viable, but future experiments or analyses could disfavor them. However, these broader issues are not very interesting for our present purposes, since all approaches to quantum gravity remain somewhat speculative. Rehashing doubts about a particular approach would not be productive unless the phenomena under consideration were to shed some specific light on these doubts. However, we do not regard the relativity of identity as evidence in its own right to doubt the hypothesis of spacetime discreteness. There is nothing about these scenarios that is contradictory or inconsistent; they are merely counterintuitive. Further, they fit naturally into the context created by previous theories as an extension of the relativity principle. We therefore focus on optimistic explanations of how such scenarios could occur and yet remain hidden.

The most obvious reason why relativity of identity might escape ordinary notice is due to the expected Planck-scale size of fundamental spacetime elements and relations involved. In simple terms, if we do not yet know how to observe *anything* about discrete spacetime structure, then we cannot expect to observe a particular discreteness-based phenomenon. However, as demonstrated by Theorem 1, relativity of identity can involve not just individual nodes and relations, but entire subgraphs of causal graphs, which could be of arbitrary size and complexity. We are reminded of Schrödinger's cat: while obvious examples of the phenomenon in question are many orders of magnitude removed from ordinary experience, no known *theoretical* reason rules out their appearance on macroscopic or even cosmological scales. We therefore turn to a second reason why such scenarios might remain hidden: entropic considerations might render them prohibitively unlikely at large scales. Here, we have occasion to revisit the notion of distinguishability that contributes to rendering these scenarios so striking. Distinguishability is typically regarded as an open-and-shut property, e.g., in sampling problems involving objects such as colored balls. However, the scenarios considered here clearly demonstrate that *mere mathematical distinguishability* can fail to supply the type of objective, robust, and permanent identity we typically associate with macroscopic objects. For example, *real* cats and dogs possess many tens of orders of magnitude of degrees of freedom, which renders their identities quite stable for practical purposes. Alternatives among small numbers of future events can conceivably swap or mix the identities of existing components of a causal structure only if those components are already quite similar or else very simple.

A third mechanism that could potentially mask the effect of relativity of identity is the potential symmetry-breaking effect of matter modeled in terms of additional fields attached to an underlying causal structure. For example, if we were to attach a "cat-valued field" to the first-generation node of the smaller connected component in the graph G in Figure 5, and a "dog-valued field" to the appropriate first-generation node of the larger connected component, then the corresponding "decorated" versions of H and H' would no longer be isomorphic, since the nodes v_1 and v_1' would possess different field values. Referring back to Section 2.1, this suggests that relativity of identity would be much more prevalent under the strong form of the causal metric hypothesis, in which all of physics is modeled in terms of causal structure, since in this case, no additional structure would be available to break potential morphisms between pairs of histories. At the opposite extreme, adding *continuous* field values to a causal structure might extinguish the possibility of relativity of identity entirely, since without further restrictions, different field values would be prohibitively unlikely to agree. In fact, such structures would likely render *all* discrete spacetime symmetry considerations emergent, since every classical history would be rigid in an absolute sense. Though it would seem ironic to model apparently smooth spacetime with discrete structure while modeling apparently granular matter-energy with continuous structure, stranger models have been proposed.

It is interesting to consider the questions of whether or not an arbitrarily advanced civilization could leverage the relativity of identity for technological purposes, and what practical use could be made of such scenarios if they could be generated. This leads to the related consideration of how such scenarios could manifest themselves to *individual* observers. A fundamental feature of the relativity principle in general seems to be that individual observers always encounter reasonable and consistent phenomena; and counterintuitive results always seem to arise from comparing the perspectives of multiple observers and attempting to reconcile them. Returning to our macroscopic cat/dog analogy, the fugue-like mixing of identities in this particular scenario seems apparent only to a "superobserver" represented by the reader, who is able to monitor a *pair* of histories involving different alternatives beginning from a common starting point. An interesting next step in exploring the relativity of identity would be to undertake a thorough analysis of which aspects of these scenarios could actually be detected by observers embedded in discrete causal histories. In such a study, observers would be modeled as nontrivial aspects of causal structure themselves, not merely as individual nodes.

Author Contributions: Conceptualization, B.F.D. and J.T.D.; methodology, B.F.D., J.T.D. and N.G.; validation, B.F.D. and N.G.; formal analysis, B.F.D. and J.T.D.; investigation, B.F.D. and J.T.D.; resources, B.F.D. and N.G.; writing—original draft preparation, B.F.D.; writing—review and editing, B.F.D. and N.G.; visualization, B.F.D. and J.T.D.; supervision, B.F.D.; project administration, B.F.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: This research does not involve external data.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bombelli, L.; Lee, J.; Meyer, D.; Sorkin, R. Space-Time as a Causal Set. *Phys. Rev. Lett.* **1987**, *59*, 5. [CrossRef] [PubMed]
2. Sorkin, R. Causal Sets: Discrete Gravity. In *Lectures on Quantum Gravity*; Gomberoff, A., Marolf, D., Eds.; Springer: New York, NY, USA, 2005.
3. Rideout, D.; Sorkin, R. Classical sequential growth dynamics for causal sets. *Phys. Rev. D* **2000**, *61*, 024002. [CrossRef]
4. Ambjorn, J.; Dasgupta, A.; Jurkiewicz, J.; Loll, R. A Lorentzian cure for Euclidean troubles. *Nucl. Phys. B Proc. Suppl.* **2002**, *106–107*, 977–979. [CrossRef]
5. Ambjorn, J.; Jordan, S.; Jurkiewica, J.; Loll, R. Second- and first-order phase transitions in causal dynamical triangulations. *Phys. Rev. D* **2012**, *85*, 124044. [CrossRef]
6. Ambjorn, J.; Jurkiewicz, J.; Loll, R. Spectral Dimension of the Universe. *Phys. Rev. Lett.* **2005**, *95*, 171301. [CrossRef]
7. Finkelstein, D. "Superconducting" Causal Nets. *Int. J. Theor. Phys.* **1988**, *27*, 473–519. [CrossRef]
8. Konopka, T.; Markopoulou, F.; Severini, S. Quantum graphity: A model of emergent locality. *Phys. Rev. D* **2008**, *77*, 104029. [CrossRef]
9. Martin, K.; Panangaden, P. A Domain of Spacetime Intervals in General Relativity. *Commun. Math. Phys.* **2006**, *267*, 563–586. [CrossRef]
10. Dribus, B.F. *Discrete Causal Theory: Emergent Spacetime and the Causal Metric Hypothesis*; Springer: Cham, Switzerland, 2017.
11. Becker, K.; Becker, M.; Schwarz, J.H. *String Theory and M-Theory*; Cambridge University Press: New York, NY, USA, 2007.
12. Rovelli, C. *Quantum Gravity*; Cambridge Monographs on Mathematical Physics; Cambridge University Press: New York, NY, USA, 2004.
13. Feynman, R. Space-Time Approach to Non-Relativistic Quantum Mechanics. *Rev. Mod. Phys.* **1948**, *20*, 367–387. [CrossRef]
14. Hawkins, E.; Markopoulou, F.; Sahlmann, H. Evolution in Quantum Causal Histories. *Class. Quantum Gravity* **2003**, *20*, 3839. [CrossRef]
15. Isham, C. Quantising on a Category. *Found. Phys.* **2005**, *35*, 271–297. [CrossRef]
16. Sorkin, R. Toward a Fundamental Theorem of Quantal Measure Theory. *Math. Struct. Comput. Sci.* **2012**, *22*, 816–852. [CrossRef]
17. Dribus, B.F. On the Axioms of Causal Set Theory. *arXiv* **2013**, arXiv:1311.2148.
18. Dribus, B.F. Entropic Phase Maps in Discrete Quantum Gravity. *Entropy* **2017**, *19*, 322. [CrossRef]
19. Hawking, S.W.; King, A.R.; McCarthy, P.J. A new topology for curved space-time which incorporates the causal, differential, and conformal structures. *J. Math. Phys.* **1976**, *17*, 174–181. [CrossRef]
20. Malament, D.B. The class of continuous timelike curves determines the topology of spacetime. *J. Math. Phys.* **1977**, *18*, 1399–1404. [CrossRef]
21. Finkelstein, D. Space-Time Code. *Phys. Rev.* **1969**, *184*, 1261–1271. [CrossRef]
22. Myrheim, J. *Statistical Geometry*; CERN Preprint; CERN: Geneva, Switzerland, 1978.

23. 't Hooft, G. Quantum Gravity: A Fundamental Problem and some Radical Ideas. In *Recent Developments in Gravitation*; Proceedings of the 1978 Cargese Summer Institute; Levy, M., Deser, S., Eds.; Springer: Boston, MA, USA, 1979.
24. Dribus, B.F. On the Foundational Assumptions of Modern Physics. In *Questioning the Foundations of Physics*; The Frontiers Collection; Springer: Cham, Switzerland, 2015; pp. 45–60.
25. Kerr, R. Gravitational field of a spinning mass as an example of algebraically special metrics. *Phys. Rev. Lett.* **1963**, *11*, 237–238. [CrossRef]
26. Bombelli, L.; Henson, J.; Sorkin, R. Discreteness without symmetry breaking: A theorem. *Mod. Phys. Lett. A* **2009**, *24*, 2579–2587. [CrossRef]
27. Sorkin, R. Does locality fail at intermediate length scales? In *Approaches to Quantum Gravity*; Oriti, D., Ed.; Cambridge University Press: New York, NY, USA, 2009.
28. Hossenfelder, S. Phenomenology of Space-time Imperfection II: Local Defects. *Phys. Rev. D* **2013**, *88*, 124031. [CrossRef]
29. Babai, L. Graph Isomorphism in Quasipolynomial Time. *arXiv* **2016**, arXiv:1512.03547.
30. Bender, E.A.; Robinson, R.W. The Asymptotic Number of Acyclic Digraphs II. *J. Comb. Theory Ser. B* **1988**, **44**, 363–369. [CrossRef]