To appear in *Annals of Mathematics and Artificial Intelligence*

The Logic of Events[1]

Glenn Shafer[2]
gshafer@andromeda.rutgers.edu

Peter R. Gillett[3]
gillett@business.rutgers.edu

Richard B. Scherl[4]
scherl@cs.njit.edu

June 18, 2000

Abstract

An event space is a set of instantaneous events that vary both in time and specificity. The concept of an event space provides a foundation for a logical—i.e., modular and open—approach to causal reasoning. In this article, we propose intuitively transparent axioms for event spaces. These axioms are constructive in the intuitionistic sense, and hence they can be used directly for causal reasoning in any computational logical framework that accommodates type theory. We also put the axioms in classical form and show that in this form they are adequate for the representation in terms of event trees established by Shafer (1998a) using stronger axioms.

---

[1] Research for this article has benefited from NSF Grant SES-9819116 to the three authors and from support for Glenn Shafer by the Fields Institute for Research in the Mathematical Sciences. We have also benefited from comments by Tatsiana Levina, Ross Shachter, and two referees.

[2] Glenn Shafer is in the Department of Accounting and Information Systems, Rutgers University, 180 University Avenue; Newark, New Jersey 07102-1897. Phone 973-353-1604. He is also Visiting Professor, Department of Computer Science, Royal Holloway College, University of London.

[3] Peter R. Gillett is in the Department of Accounting and Information Systems, Rutgers University, Janice H. Levin Building, 94 Rockafeller Road; Piscataway, New Jersey 08854-8054. Phone 732-445-4765.

[4] Richard B. Scherl is in the Department of Computer and Information Science, New Jersey Institute of Technology, Newark, New Jersey 07102-1982. Phone 973-596-2657.

Table of Contents

I. Introduction

An event space is a set of instantaneous events that vary both in time and specificity. The concept of an event space provides a foundation for a logical—i.e., modular and open—approach to causal reasoning. Mathematically, event spaces generalize both event trees and Boolean algebras. Shafer (1998a) gave axioms for event spaces and demonstrated the adequacy of these axioms by showing that they lead to a representation in terms of an event tree, generalizing Stone's representation of a Boolean algebra in terms of a field of subsets. Unfortunately, Shafer's axioms are not intuitively transparent and do not lend themselves to computer implementation. They are also unnecessarily strong in some respects. In particular, they make the assumption that the failure of an event is itself a well-specified event.

In this article, we provide simpler and more transparent axioms for event spaces. These axioms are constructive in the intuitionistic sense, which means that they can be used directly for causal reasoning in any computational system that supports higher-level type theory. We analyze the axioms carefully from a constructive point of view—i.e., without the principle of the excluded middle. We also translate them into classical form and show that in this form they are adequate for the representation proven by Shafer (1998a) with his stronger axioms.

Because the concept of an event space is not yet widely understood, we preface our axiomatization with an extended intuitive explanation of how the concept arises from the study of event trees. This explanation, in Part II, includes a discussion of how the concept of an event space can be generalized to accommodate the theory of relativity. Part III, which formulates and studies our constructive axioms, is the heart of the article. In Part IV we translate the axioms into classical form and prove the representation theorem.

In Part V, we discuss briefly how our axioms can be used directly as a logic in higher-level type theory, and in Part VI we compare our framework with other approaches to temporal and causal reasoning. Some extensions are discussed in an appendix.

II. An Informal Look at Event Spaces

Event spaces embed Boolean algebras in event trees. When we think of the elements of a Boolean algebra as events, these events are ordered by specificity; $E \subseteq F$ means that the event E narrows the event F by specifying more detail about what happens. In an event tree, on the other hand, the partial order is temporal: $E \le F$ means that the event E happens after the event F.[5] In an event space, these two partial orders coexist.

As we will now explain in detail, the events in an event space can be represented as sets of nodes in an event tree. This representation makes clear how the partial orders $\subseteq$ and $\le$ can coexist: $E \subseteq F$ means that the set representing E is contained in the set representing F, while $E \le F$ means that the set representing E comes later in the tree than the set representing F.

We begin, in §§1-2, by reviewing the concept of an event tree mathematically and philosophically. In §3, we discuss when a set of nodes in an event tree can represent an instantaneous event (no node in the set can lie below another one in the tree). In §§4-6, we discuss our two partial orders: $\subseteq$, under which the instantaneous events almost form a Boolean algebra, and $\le$, under which they form a distributive lattice. In §§7-8, we discuss additional constructions that were emphasized by Shafer (1998a). Finally, in §9, we discuss how the concept of an event space can be generalized to accommodate the theory of relativity.

This part of the article is heuristic and intuitive. Intuitively, the events in event spaces correspond to sets of nodes in event trees, and because event trees are relatively well understood, we can use this correspondence to develop our intuitions about event spaces. The rest of the article will be more formal. In Parts III and IV, we abstract formal axioms for event spaces from the intuitions developed here. In Part V, we show that these axioms are faithful to the intuitions by showing that a space satisfying them is isomorphic to a space of sets of nodes in an event tree.

1. Event Trees

Mathematically a *tree* is a set with a partial order $\le^t$ in which any two elements with a lower bound are comparable ($G \le^t E$ and $G \le^t F$ together imply that $E \le^t F$ or $F \le^t E$). This idea is illustrated in Figure 1.

---

[5] It is matter of convention whether we write $E \le F$ or $F \le E$ when E happens after F. Because we are accustomed to representing earlier times with smaller numbers, it may seem natural to put the earlier event on the left. We choose instead to put the earlier event on the right because, under the precise definitions that we will adopt, this makes $E \subseteq F$ a special case of $E \le F$. The alternative, which is unnecessarily confusing, would be for $E \subseteq F$ to be a special case of $F \le E$.

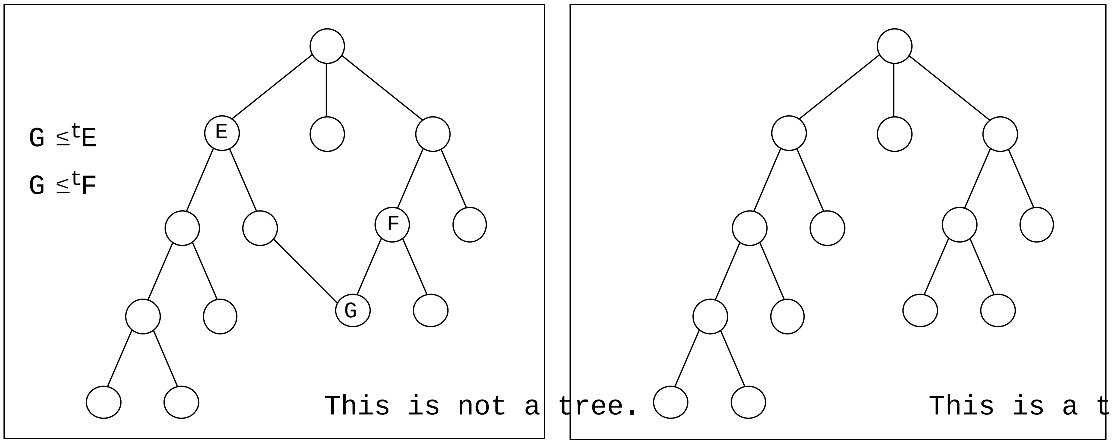

Figure 1 When a partially ordered set is finite, we can arrange its elements in a diagram in which E $\leq^t$ F if and only if there is a path downward from F to E. The partially ordered set depicted on the left is not a tree, because the elements E and F have an element G below them both even though they are not comparable in the partial ordering (neither is on a path below the other). The partially ordered set depicted on the right, in contrast, is a tree; it has no incomparable pair E and F with an element G below both.

Calling a tree an *event tree* means that we interpret its elements as instantaneous events and that we interpret the partial order as temporal order; E $\leq^t$ F means that E, if it happens, happens at the same time as F or after F. This permits E and F to be equal. If E $\leq^t$ F but E is not equal to F, then E can only happen if F happens strictly earlier; in this case we may write E $<^t$ F.

An event tree is not necessarily finite, but we will illustrate our ideas using finite diagrams as in Figures 1 and 2. These diagrams will follow the convention (common in the statistics and economics literature but opposite to the most common convention in philosophy and physics) that time runs downward. A path downward represents a *history*—one possibility for how events may evolve.

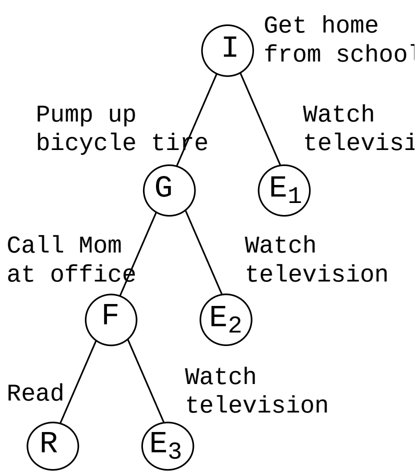

Figure 2 An event tree for what Rick may do after school. He may watch television right away, he may delay watching television, or he may even end up reading instead of watching television. We assume that the different paths down the tree represent all the ways in which the events shown can happen. At the outset, for some reason, we can rule out the possibility that Rick might call his mother and then pump up his bicycle tire afterwards.

A node E in an event tree has a dual meaning. On the one hand, E is an instantaneous event—something that happens at a particular instant. On the other hand, E is a situation that the world is in at a particular instant. The situation E is the situation that arises when the event E happens. The event E is the event that the situation E arises. The node F in Figure 2, for example, represents both the event that Rick calls his mother and the situation in which he does so.

The interpretation of a tree as an event tree involves a number of assumptions, which are not always made in other contexts where the words "situation" or "event" are used, and which should therefore be stated explicitly. These assumptions are explained most easily in the finite case, where we can talk about the daughters of a situation—the situations immediately below it.[6]

1. A daughter cannot happen until after its mother happens (strictly after; not at the same time).
2. Two distinct daughters of a situation are mutually exclusive; they cannot both happen. In situation I in Figure 2, Rick may pump up his bicycle tire (G) or watch television right away (E_1), but he cannot do both.
3. The daughters of a situation (and even all paths down from the situation) are all possible in that situation, no matter what else may be said about what has happened. In situation G in Figure 2, where Rick has just pumped up his bicycle tire, it is possible that he will call his mother at the office and then read, no matter how much pressure he put in the bicycle tire.[7]
4. The daughters of a situation are exhaustive; once the world is in that situation, one of the daughters must happen. In situation I in Figure 2, if Rick does not pump up his bicycle tire (G), then he watches television right away (E_1). There is no third possibility.
5. In order to specify a situation fully, we must say how we got there. In other words, we regard situations that are similar but are preceded by different histories as different situations. This is why we assume that the situations form a tree. If we can get into a situation G after being in a situation E, then we cannot also get into G after being in a situation F that is not in the same history (path down the tree) as E. If $G \leq^t E$ and $G \leq^t F$, then $E \leq^t F$ or $F \leq^t E$.
6. By the same token, the same situation cannot arise twice as the world evolves (moves down the tree). And hence the same instantaneous event cannot happen twice.
7. On the other hand, a situation does not specify completely everything that has happened in the world. It is only as detailed in its meaning as the tree in which it is situated.

Another way of bringing out the density of meaning in an event tree is to consider the different ways of explaining what is meant by $E < F$. One way of explaining this relation is to say that E can only happen if F has already happened. Another explanation is that in F, E's later happening is possible. When E and F are represented by nodes in the same event

---

6

7 This assumption is not made in all cases. We insist on it here only because it gives us a clear causal story from which to proceed. Our purpose is not to standardize the causal assumptions made when trees are used but to develop an abstract language in which causal assumptions can be stated precisely, whether or not they can be or are represented in a tree.

tree, these two conditions, which seem to have rather different content, are equivalent. As we will see shortly, the two conditions do indeed have different content and are not equivalent in general in an event space.

2. Philosophical Aside: Nature as Witness

From a philosophical point of view, the most challenging question raised by an event tree is the meaning of possibility. When we explain Figure 2, we say that it is possible at the outset for Rick to call his mother or to watch television right away, and then we say there is no third possibility. It is impossible, for example, that he will read before he calls his mother. What does this mean and how could it be known?

Over the centuries, philosophers have expounded many divergent doctrines about the meaning of possibility. Some believe that because God knows exactly what will happen, nothing else is possible. This eliminates our whole enterprise, by reducing our event tree to a straight line showing what actually happens. Others take the view that possibility is relative to knowledge and is therefore subjective or personal. If Peter knows more than Paul, what is possible in Paul's event tree may not be possible in Peter's.

We will not try to cast new light on the meaning of possibility. Rather, we will work under the everyday assumption that possibility is objective and evolves with time. An event is at first objectively possible and then either happens or becomes impossible as the world evolves. In order to acknowledge the force of the thesis that possibility must be relative to knowledge, we will imagine a super scientist or demigod whose knowledge evolves in step with this objective possibility. Objective possibility is the same, we may suppose, as possibility for this demigod. Her knowledge surpasses that of all actual agents, human or artificial, but falls short of God's perfect knowledge. Time does not pass for God, and events do not happen, because God has already foreseen everything at the beginning of time. But our demigod's knowledge increases with time, thus defining what is possible and marking the happening of events. An event is possible if the demigod has not yet ruled it out, it happens when she witnesses it, and it becomes impossible when she rules it out.

The concept of a demigod whose superior knowledge defines objective possibility has a long history. Laplace imagined such a demigod, whom he called *l'intelligence superieure*, in order to explain determinism (Bru 1986), and Cournot used the same concept to explain objective probability (Martin 1996). Shafer (1996) calls the demigod *Nature* and puts her at the center of his philosophy of causality. From time to time in this article we will make use of *Nature* in our intuitive explanations.

3. Clades

In our first look at event trees, in § 1, we emphasized that individual nodes in an event tree can be interpreted as instantaneous events or situations. As we will now emphasize, we can also sometimes group nodes together and interpret the whole set as an instantaneous event or situation. In Figure 2, for example, the instantaneous event that Rick calls his mother is represented by the single node F, while the instantaneous event that he watches television is represented by the three nodes E_1, E_2, and E_3, taken together.

The concept of instantaneous event that we use in this article allows such an event to happen only once as the world evolves. So we will say that a set of nodes in an event tree represents an instantaneous event if and only if none of its elements precedes another on a path down the tree. We call a set of nodes satisfying this condition a *clade*.

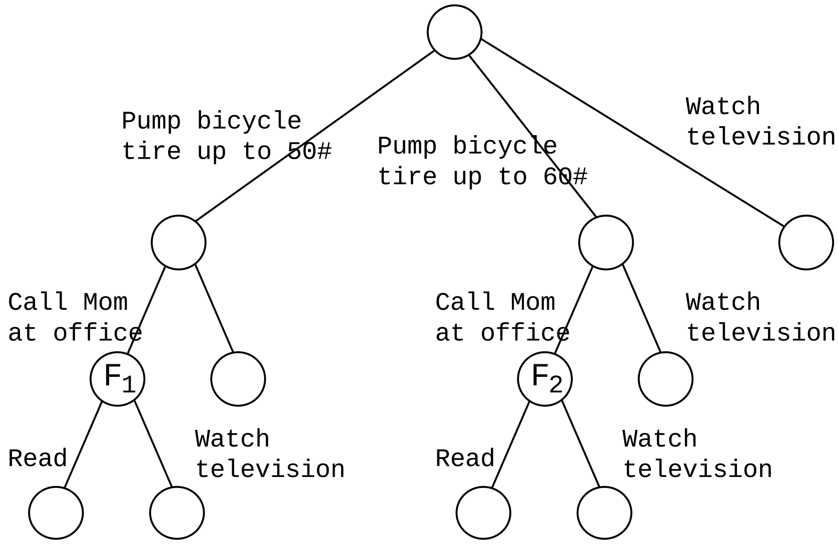

Figure 3 More or less air in the bicycle tire.

We should not think of an instantaneous event represented by a single node in a particular event tree as fundamentally different from an instantaneous event represented by larger clade, since the difference lies in our choice of representation, not in the events themselves. Indeed, any instantaneous event represented in one event tree by a clade consisting of several nodes can equally well be represented in a different event tree by a single node, and vice-versa.[8] Figures 3 and 4 illustrate the point. In Figure 3, we add detail to Figure 2, by specifying how much air Rick puts in his tire, thus turning the single node F, which represents the event that Rick calls his mother, into the clade $\{F_1,F_2\}$. In Figure 4, we go in the opposite direction, removing detail from Figure 2 and thus representing Rick's watching television, which appears as the clade $\{E_1,E_2,E_3\}$ in Figure 2, as a single node E.

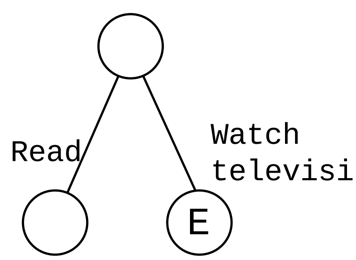

Figure 4 Watching television as a single node.

We can draw event trees that disagree. But the event trees in Figures 2, 3, and 4 do not disagree. They merely provide different levels of detail. Figure 3 says more than Figure 2,

---

[8] A note of caution: Although any single instantaneous event can be represented by a single node in some event tree (for example, an event tree with only a single node, which is taken to represent that event), there are cases where two instantaneous events cannot be represented by single nodes in the same event tree.

and Figure 4 says less, but none of the three deny any assertion about possibility or impossibility made by one of the others. For example, Figure 2 says that in the situation where Rick pumps up his bicycle tire, it is possible that he will read later, and Figure 3 agrees; it tells us that no matter which pressure he puts in the tire, it is possible that he will read later.

An event space consists of instantaneous events of the kind that can be represented as clades. But it treats these events abstractly, without reference to their representation in any particular event tree. An element of an event space will appear as a single node in some event trees, as a clade in others. As a result, the events in an event space can relate to each other in a great variety of ways. One instantaneous event may precede another (as F precedes E_3 in Figure 2), and one may refine another (as E_3 refines E in Figure 2), but often the relation is more complex. There is, for example, no simple description of the relation between {F,E_2} and E in Figure 2.

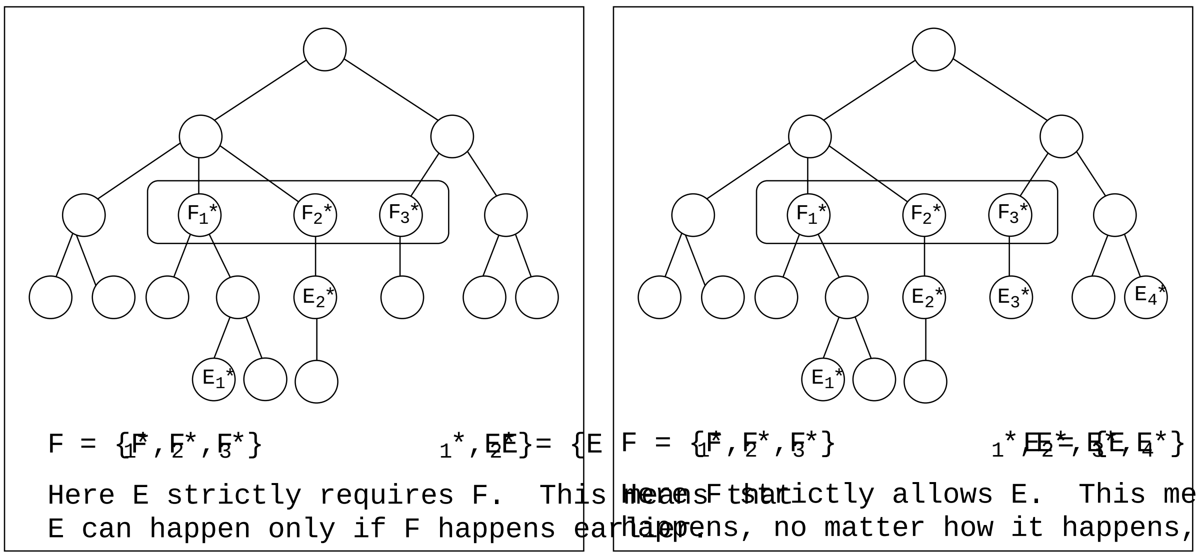

Figure 5 On the left, E requires F, but we cannot say that F allows E; there is one way F can happen without E being possible later, namely F_3*. On the right, F allows E, but we cannot say that E requires F; there is one way E can happen without F happening first, namely E_4*.

One aspect of the diversity of relations among instantaneous events is the diversity of meanings that can be given to the statement that one event comes before or after another. In the case of single nodes, there is no ambiguity. When we say that F strictly precedes E, this means E is strictly below F in the tree; there is a path from F down to E. This implies both (1) that E can happen only if F happens strictly earlier (in this case, we say that E *strictly requires* F), and (2) that E's later happening is possible in the situation where F happens, no matter what else happens then (in this case, we say that F *strictly allows* E). When we consider clades consisting of larger numbers of nodes, these two conditions are not equivalent, as Figure 5 illustrates.

4. The Refinement Order

Suppose E and F are instantaneous events, and suppose that whenever E happens, F happens simultaneously. In this case, we say that E *refines* F, and we write $E \subseteq F$. When E and F are represented as clades in the same event tree, E ⊆ F means that the set representing E is a subset of the set representing F. Since set inclusion is a partial order, ⊆ is a partial order on instantaneous events. This partial order has a zero, the impossible event, which we will designate by Λ, and which is represented in an event tree by the empty clade.[9]

Any two instantaneous events E and F have a greatest lower bound in the order ⊆, which is represented in an event tree by the their set-theoretic intersection E∩F. We call E∩F the *overlap* of E and F. Intuitively, E∩F is the event that E and F happen simultaneously. This event may be possible, as in Figures 7 and 8, or impossible, as in Figure 6.

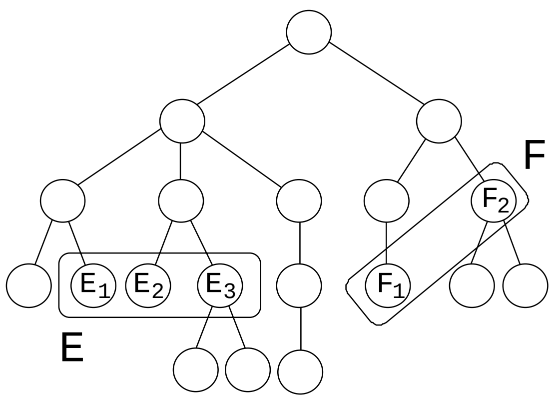

$E = \{E_1,E_2,E_3\}$
$F = \{F_1,F_2\}$

$E \cap F = \Lambda$
$E \cup F = \{E_1,E_2,E_3,F_1,F_2\}$

Figure 6 Two events that do not overlap and can be merged.

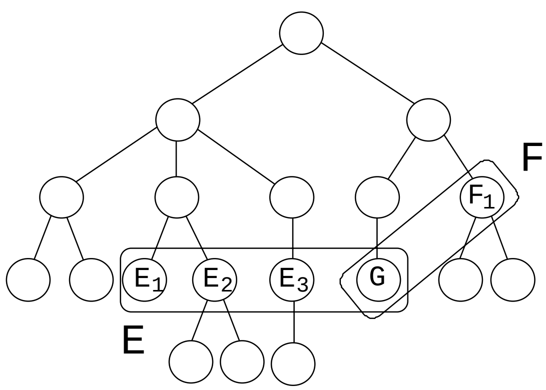

$E = \{E_1,E_2,E_3,G\}$
$F = \{G,F_1,\}$

$E \cap F = \{G\}$
$E \cup F = \{E_1,E_2,E_3,G,F_1\}$

Figure 7 Two events that overlap and can be merged.

9 The impossible event is the only instantaneous event that cannot also be thought of as a situation.

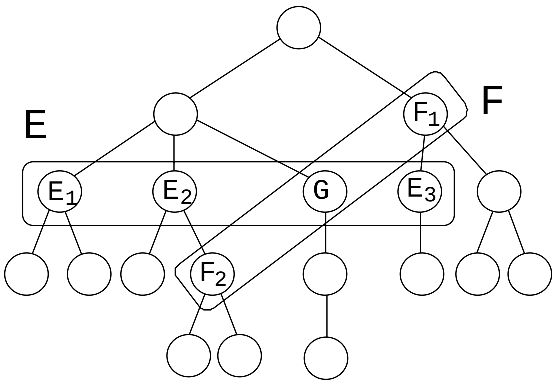

$E = \{E_1, E_2, E_3, G\}$
$F = \{F_1, F_2, G\}$

$E \cap F = \{G\}$
$E \cup F$ does not exist. (E_2 and F_2 cannot be in the same clade, and E_3 and F_1 cannot be in the same clade.)

Figure 8 Two events that overlap and cannot be merged.

The story about upper bounds is not quite so simple. As Figure 8 illustrates, two instantaneous events E and F may fail to have an upper bound in E and F. This happens when E and F are represented by clades such that a node in one is above a node in the other on some path. Because a clade cannot contain two such nodes, no clade can encompass both E and F in this case; there is no instantaneous event that happens simultaneously whenever E happens and also whenever F happens. If E and F do have an upper bound, then they have a least upper bound, which we designate by $E \cup F$ and call the *merger* of E and F.

Given two instantaneous events E and F, we can also form the event that E happens without F happening at the same time. This event, designated by $E \backslash F$, is the *complement* of F relative to E.

Except for the fact that the least upper bound of E and F does not always exist, the partial order $\subseteq$ has all the properties of a Boolean algebra: it is distributive, unique relative complements always exist, etc. In fact, if we consider only instantaneous events that refine a fixed instantaneous event I, then we do have a Boolean algebra: the Boolean algebra consisting of all subsets of I.

5. The Temporal Lattice

Suppose E and F are instantaneous events, and suppose that whenever E happens, F either happens simultaneously or else has already happened earlier. In this case, we say that E *requires* F, and we write $E \leq F$.[10] If E and F are represented as clades in the same event tree, then $E \leq F$ means that each node in E is either (1) in F or (2) on a path down the tree from a node in F. As it turns out, $\leq$ is also a partial order on instantaneous events. It has the same zero as $\subseteq$: the impossible event.

It is clear from the definition of $\leq$ that if $E \subseteq F$, then $E \leq F$. Another special case of $E \leq F$ is the case illustrated on the left of Figure 5—E strictly requires F. When E refines F ($E \subseteq F$), all of E is inside F. When E strictly requires F, none of E is inside F. These are two extremes; in the general case where $E \leq F$, some of E may be inside F while the remainder is below F. Notice that "E strictly requires F" is a stronger condition than "$E \leq F$ but $E \neq F$."

---

10 This is the same as the relation on the left of Figure 5, except that E and F may overlap.

This is another aspect of how clades in an event tree relate to each other in more complicated ways than individual nodes.

Given any two instantaneous events E and F, we may define a new event: the event that E happens after F. (Here we use "after" in a broad sense, to mean after or at the same time). We designate this event E^F. In an event tree in which both E and F appear as clades, E^F will be the clade consisting of all nodes in E that are also in F or else on a path down from a node in F. Notice that $E \leq F$ if and only if $E = E^F$. The construction E^F is illustrated in Figure 9.

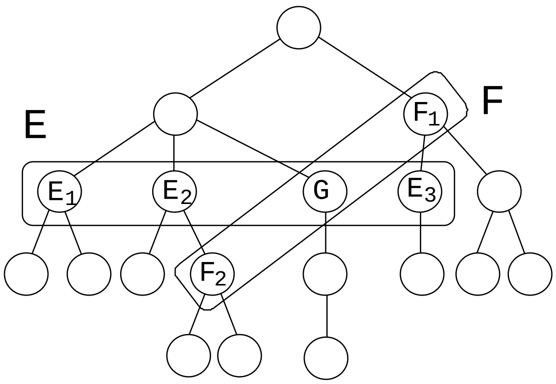

$E = \{E_1,E_2,E_3,G\}$
$F = \{F_1,F_2,G\}$

$E^F = \{G,E_3\}$
$F^E = \{G,F_2\}$
When G happens, E and F happen at the same time. When E_3 happens, E happens strictly after F. When F_2 happens, F happens strictly after E.

Figure 9 One event happening after another.

The significance for causal thinking of the temporal order ≤ and the construction E^F may not be obvious at first glance, but we can use them to define relations that clearly do have causal meaning. Here are four examples:

- If $G \leq E$ and $G \leq F$ imply $G = \Lambda$, then we say E and F *diverge*, and we write div(E,F). This means that they cannot both happen (because there is no situation H where they both have happened). This relation is symmetric; div(E,F) if and only if div(F,E). But we can rephrase it in a way that seems asymmetric: the happening of E makes F impossible, if F was already not impossible. Or we can say that in the situation E, F is impossible. Figure 6 gives an example where E and F diverge.
- If div(G,F) implies div(G,E), then we say E *implies* F, and we write $E \rightarrow F$. This means that whenever happening of F is ruled out (because we are in situation G that diverges from F), the happening of E is also ruled out. So if E happens, and hence will never be ruled out, F also will never be ruled out and hence must happen at some point. (We assume that every event eventually happens or fails.) Figure 10 gives an example.
- If $G \subseteq E$ and $G \neq \Lambda$ imply $F^G \neq \Lambda$, then we say E *allows* F, and we write E ◊→ F. This means that in the situation E, F is possible, no matter what else has happened.
- When $E \leq F$ and F ◊→ E, we say F *precedes* E. This has the same meaning as the partial order in an event tree. Two instantaneous events E and F can be represented in an event tree where they satisfy $E \leq^t F$ if and only if F precedes E. If one of the two relations $E \leq F$ and F ◊→ E is satisfied but the other is not (as in Figure 5), then there is no event tree in which both events can be represented. See (Shafer 1998a) for further discussion.

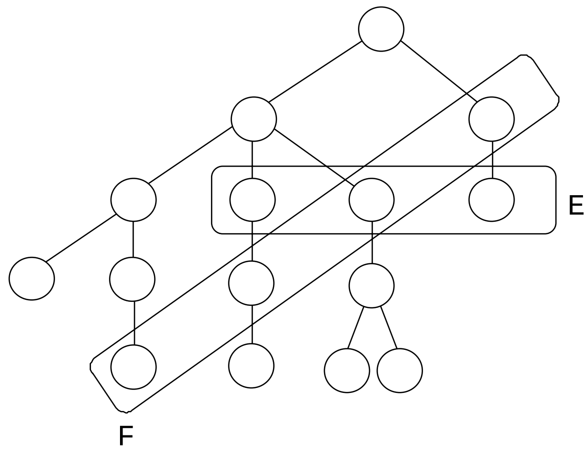

Figure 10 E implies F.

As it turns out, instantaneous events form a distributive lattice with respect to the temporal order ≤. This means that every two instantaneous events E and F have a greatest lower bound E∧F and a least upper bound E∨F, and ∧ and ∨ obey the distributive laws. We call E∧F the *ending* of E and F. It is the instantaneous event that E and F finish happening, and it can happen in three different ways: E and F may happen at the same time, E may happen after F has already happened, or F may happen after E has already happened. We call E∨F the *beginning* of E and F. It is the instantaneous event that E and F begin to happen, by at least one of them happening. It also can happen in three different ways: E and F may happen at the same time, E may happen without F yet having happened, or F may happen with E yet having happened.

The ending and beginning of E and F can be defined in terms of the constructions we have already mentioned: Overlap (E∩F), Merger (E∪F), Complement (E\F), and After (E^F). Indeed,

$$E \wedge F = E^F \cup F^E, \tag{II.5.1}$$

and

$$E \vee F = (E \backslash (E^F)) \cup (F \backslash (F^E)) \cup (E \cap F). \tag{II.5.2}$$

Because E∩F is a refinement of both E^F and F^E, it is contained in E∧F as well as in E∨F.

In Figure 8, the beginning and ending are given by E∧F = $\{F_2, G, E_3\}$ and E∨F = $\{E_1, E_2, G, F_1\}$. The configuration of these events in the event tree is depicted schematically in Figure 11.

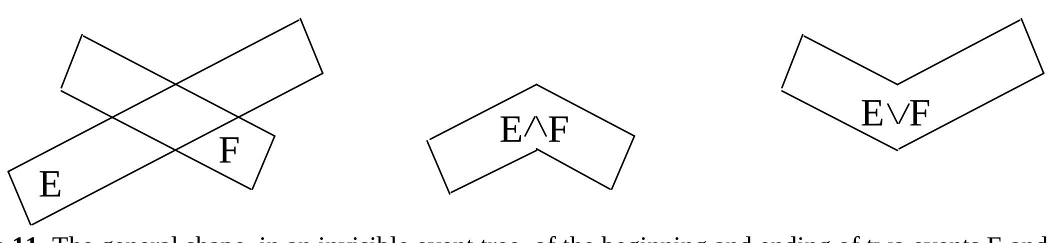

Figure 11 The general shape, in an invisible event tree, of the beginning and ending of two events E and F. Here, as in Figure 8, we suppose that the two events lie across each other in the event tree, so that either can happen before the other. The ending is then the lower bow, possibly with its wings chopped off, while the beginning is the upper bow, with the full extent of its wings.

If the merger $E \cup F$ exists, then it is the same as $E \vee F$, and then $E \cap F$ is also the same as $E \wedge F$. This happens in Figures 6 and 7.

6. A Simple Temporal Logic

As we shall see in Parts II and III, we can provide a full set of axioms for event spaces using only the two partial orders $\subseteq$ and $\leq$ and the five fundamental constructions $E \cap F$, $E \cup F$, $E \backslash F$, Λ, and E^F. The axioms for these two partial orders and five constructions constitute a simple but surprisingly powerful and flexible temporal and causal logic.

One aspect of the power of this logic is its ability to encode the information in an event tree. Figure 12 shows a straightforward encoding for a simple tree consisting of a mother and daughters. By repeating this encoding for each mother, we can encode the information contained in any finite event tree.

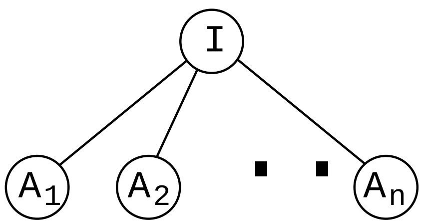

1. $A_i \leq I$ and $A_i \cap I = \Lambda$, for $i = 1,\ldots,n$.
2. $I \Diamond\!\!\rightarrow A_i$, for $i = 1,\ldots,n$.
3. $\mathrm{div}(A_i,A_j)$, $1 \leq i,j \leq n$.
4. $I \rightarrow A_1 \cup A_2 \cup \cdots \cup A_n$.

Figure 12 The information in an event tree. The four conditions correspond to the first four points about the interpretation of an event tree in §1. They do not include any requirement that the A_I be "immediate" successors of I. In an event space, we always leave open the possibility that additional situations may be interpolated between a given situation and a later situation.

More importantly, our logic can encode information that does not fit the mold provided by the concept of an event tree. We often find ourselves working with events $I,A_1,A_2,\ldots,A_n$ about which we know some but not all of the information listed in Figure 12. When this is the case, the modularity of our logic permits us to state the information we do have and draw inferences from it. Research workers in disciplines such as operations research and decision analysis, which use models that incorporate event trees, often note that the most difficult part of their work is the modeling step. The analyst must put great effort into finding or conjecturing enough information to define a tree, even when this information is not logically relevant to the inferences that are needed. A more modular approach obviously allows the analyst to concentrate instead on information that is really relevant.

Even when the available information is equivalent to that in a tree, it may be more natural to elicit and express it logically. Consider Figure 2, where we began with these events:

I = Get home from school.
G = Pump up bicycle tire.
E = Watch television.
F = Call Mom at office.
R = Read.

Figure 2 suggests that we express our information about E in terms of the decomposition E_1, E_2, and E_3, where

$$E_1 = E \backslash E^G, \qquad E_2 = E^G \backslash E^F, \qquad E_3 = E^F.$$

This leads to the logical statements on the left in the following table.

Figure 2's information expressed in terms of I, G, F, R, E_1, E_2, and E_3.	**Figure 2's information expressed in terms of I, G, E, F, and R.**
$G \leq I$, $G \cap I = \Lambda$, $E_1 \leq I$, $E_1 \cap I = \Lambda$ $F \leq G$, $F \cap G = \Lambda$, $E_2 \leq G$, $E_2 \cap G = \Lambda$ $R \leq F$, $R \cap F = \Lambda$, $E_3 \leq F$, $E_3 \cap F = \Lambda$	$G \leq I$, $G \cap I = \Lambda$, $E \leq I$, $I^{E} = \Lambda$ $F \leq G$, $F \cap G = \Lambda$, $G^{E} = \Lambda$ $R \leq F$, $R \cap F = \Lambda$, $F^{E} = \Lambda$
$I \Diamond\!\!\rightarrow G$, $I \Diamond\!\!\rightarrow E_1$ $G \Diamond\!\!\rightarrow F$, $G \Diamond\!\!\rightarrow E_2$ $F \Diamond\!\!\rightarrow R$, $F \Diamond\!\!\rightarrow E_3$	$I \Diamond\!\!\rightarrow G$, $I \Diamond\!\!\rightarrow E$ $G \Diamond\!\!\rightarrow F$, $G \Diamond\!\!\rightarrow E$ $F \Diamond\!\!\rightarrow R$, $F \Diamond\!\!\rightarrow E$
div(G,E_1) div(F,E_2) div(R,E_3)	div(R,E)
$I \rightarrow G \cup E_1$ $G \rightarrow F \cup E_2$ $F \rightarrow R \cup E_3$	$I \rightarrow G \vee E$ $G \rightarrow F \vee E$ $F \rightarrow R \vee E$

It is more likely, however, that our knowledge begins as knowledge about E, and it is clearly simpler to express it in this form, as in the right column of the table. (We will leave it to the reader to verify, using the axioms in Part IV, that the statements on the right are equivalent, taken together, to those on the left.)

The advantages of a logical notation over a graphical representation grow as we add more events to the conversation. Even if the new events involve little additional information, they may require a substantially enlarged event tree, with great repetition. This can be illustrated by the progression from Figure 4, where we have only the three events I, R, and E, to Figure 2, where the events F and G are interpolated, to Figure 3, where G is decomposed into two parts, G_1 and G_2. As we see in the last panel of the following table, the decomposition of G involves only four simple statements in the logic, whereas in Figure 3 it involves the duplication of an entire branch of the tree.

Graphical Representation	**Logical Representation**
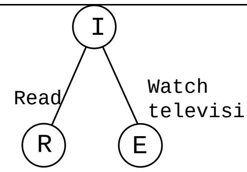	$R \le I$, $E \le I$, $R \cap I = \Lambda$, $E \cap I = \Lambda$ $I \Diamond\!\!\rightarrow R$, $I \Diamond\!\!\rightarrow E$ div(R,E) $I \rightarrow R \cup E$
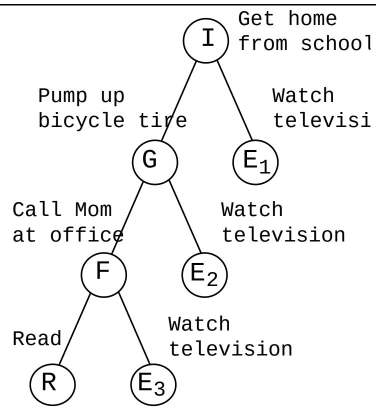	Add to preceding: $G \le I$, $G \cap I = \Lambda$ $F \le G$, $F \cap G = \Lambda$, $G^E = \Lambda$ $R \le F$, $R \cap F = \Lambda$, $F^E = \Lambda$ $I \Diamond\!\!\rightarrow G$ $G \Diamond\!\!\rightarrow F$, $G \Diamond\!\!\rightarrow E$, $F \Diamond\!\!\rightarrow R$, $F \Diamond\!\!\rightarrow E$ div(R,E) $I \rightarrow G \vee E$, $G \rightarrow F \vee E$, $F \rightarrow R \vee E$
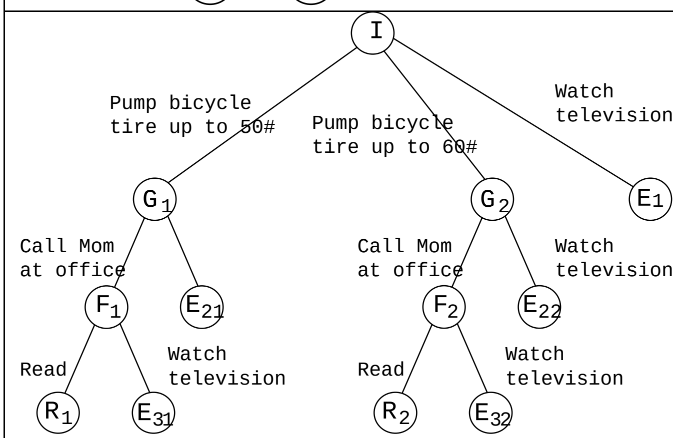	Add to preceding: $G_1 \ne \Lambda$ $G_2 \ne \Lambda$ $G_1 \cap G_2 = \Lambda$ $G = G_1 \cup G_2$

7. Generalizing to a Relativistic World[11]

This article is motivated by a desire to develop systems for ordinary causal reasoning, not reasoning systems for space travelers. A thorough logical analysis of temporal relations should, nevertheless, respect the insights of Einstein's theory of relativity. In the world as Einstein has taught us to understand it, events do not all lie along a single universal time line. An event E can be said to precede an event F only if the locations of E and F in space and time permit the news of E's happening, traveling at the speed of light from where E happens to where F happens, to arrive by the time F happens. If the two events are far apart in space, it may be that neither precedes the other in this sense. They are incomparable. Yet if they both happen there is a later situation where they are both in the past—i.e., where the news of both has arrived.

It is very easy to see that the concept of an event tree does not accommodate Einstein's insights. Indeed, his fundamental insight is violated by the condition that makes a partially

[11] The ideas in this section are not needed for an understanding of the remainder of the article, but they motivate the organization of the last few sections of Part III.

ordered set a tree—the condition that if both E and F are in the past of G ($G \leq E$ and $G \leq F$), then either E or F precedes the other ($E \leq F$ or $F \leq E$).

From a philosophical point of view, the difficulty can be located in the fact that an event tree represents the possibilities for the increasing knowledge of a single witness. As we explained in §2, this witness, Nature, always knows everything that has happened so far. Her knowledge grows only inasmuch as she learns new things as they happen. A relativistic generalization of this picture would require a plethora of witnesses, who follow different paths through space-time. In a given situation E, we will find many of these witnesses, who have followed different paths to get to the same place at the same time. In E they will all know the same things; for they will have pooled their knowledge, so to speak, thus coming to know everything that has happened in the Minkowski cone that constitutes the past of that situation. But as one of these witnesses moves forward in time, her past Minkowski cone grows; she learns not only about what happens as she moves forward but also more about what has already happened. Events enter her past without passing through her present.

The axiomatization of event spaces we provide in Part III is faithful to the picture of an event tree and hence is inconsistent, as a whole, with relativity. But it can easily be generalized so it is consistent with relativity. The relativistic generalization retains the refinement order $\subseteq$, the temporal order $\leq$, and the five fundamental constructions $E \cap F$, $E \cup F$, $E \backslash F$, Λ, and E^F, with some clarifications in their meaning. We lose only two final axioms, which reduce the structure to a tree and permit the construction of the ending $E \wedge F$ and the beginning $E \vee F$.

The first step in the relativistic generalization is to clarify the meaning of the refinement order $\subseteq$. We have said that $E \subseteq F$ means that whenever E happens, F happens simultaneously. In a non-relativistic world, "simultaneous" means simply "at the same time." But in a relativistic world, it means "at the same time and at the same place." This must be kept in mind in interpreting all the related constructions. For example, the greatest lower bound $E \cap F$ is characterized by two axioms, which we can state here in this form:[12]

Axiom 2A $E \cap F \subseteq E$ and $E \cap F \subseteq F$.

Axiom 2B If $G \subseteq E$ and $G \subseteq F$, then $G \subseteq E \cap F$.

According to these axioms, $E \cap F$ is the event that E and F happen at the same time and the same place. (Axiom 2A says that when $E \cap F$ happens, both E and F happen at the same time and same place, and Axiom 2A says that when E and F both happen at the same time and place, $E \cap F$ happens there, too.)

The least upper bound $E \cup F$ does not always exist, but the condition for its existence and its characterization when it does exist can be stated in this way:[13]

Axiom 3 If $E \subseteq I$ and $F \subseteq I$, then $E \cup F$ can be formed, and $E \cup F$ then satisfies $E \subseteq E \cup F$, and $F \subseteq E \cup F$, and $E \cup F \subseteq I$.

This says that $E \cup F$ happens at the same time and place as E whenever E happens and at the same time and place as F whenever F happens. It also has an interesting further consequence: if E and F are both refinements of the same event I, and they both happen, then they must happen at the same time and place as each other (because they both happen at the same time and place as $E \cup F$ and I).

[12] See §2 of Part III.

[13] See §3 of Part III.

The condition that two refinements of an instantaneous event must happen at the same time and place if they both happen constitutes a restriction on what we are allowed to classify as an instantaneous event. Compare, for example, these two purported instantaneous events:

- G = "Joe dies."
- H = "Joe dies on earth during Year A, or Bill dies on a planet around Sirius during year B," where Year A on earth and Year B on Sirius are at a space-like distance—light leaving earth during Year A will not reach Sirius before the end of Year B there, and vice versa.

The event G qualifies as an instantaneous event in our system, even if Joe is a space traveler. There are many refinements of G—many times, places, and ways Joe can die. But if Joe dies in way 1 and Joe dies in way 2, then these two events happen at the same time and the same place. On the other hand, H clearly does not qualify as an instantaneous event, for then Joe dying on earth during Year A and Bill dying on Sirius during Year B would both be refinements of H, these events, which can both happen, obviously cannot happen in the same time and place.

We must also clarify the meaning of the temporal order $\leq$. In the non-relativistic case, the statement $E \leq F$ means that whenever E happens, no matter how it happens, F has already happened. We can use these same words in the relativistic generalization, provided we are clear about what "F has already happened" means; it means that the happening of F is the past Minkowski cone: F has happened at a place and time such that the news of its happening, traveling at the speed of light, has had time to arrive when and where E happens. With this interpretation, most of the axioms for $\leq$ and for the construction E^F appear to be valid for the relativistic case.

The conflict with relativity theory is limited to the two final axioms in our system, which can be paraphrased as follows:[14]

Axiom 14 If neither E nor F can happen after the other, the merger $E \cup F$ can be formed.

Axiom 15 $G \leq E$ and $G \leq F$ imply $G \leq E^F \cup F^E$.

Fortunately, a great deal of ordinary causal reasoning can be accomplished without these two axioms.

The difficulty with Axioms 14 and 15 is illustrated in Figure 13, which depicts the possible trajectories of two witnesses, Solid and Dashed. Solid and Dashed are together at two points in the story, first in situation I and then, depending on how events turn out, in one of four possible later situations, A, B, C, or D. In the interim, they travel apart and witness different events. For simplicity, we suppose that they travel at the speed of light. Solid, whose event tree is drawn with solid lines, witnesses either E or H, while Dashed, whose event tree is drawn with dashed lines, witnesses either F or G. Because they are traveling at the speed of light, it is only when they reunite, in A, B, C, or D, that each learns what the other witnessed. If Solid witnessed E and Dashed witnessed F, then they reunite in A, if Solid witnessed E and Dashed witnessed G, then they reunite in B, etc.

[14] See §§14-15 of Part III.

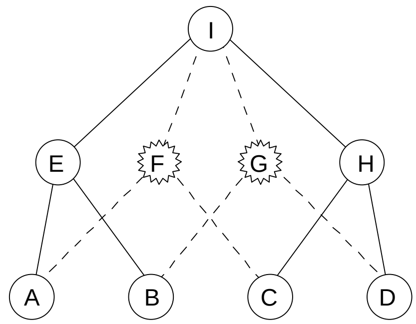

Figure 13 Intertwined event trees for two witnesses, Solid and Dashed. They begin together in situation I, go their separate ways to witness different events, and then meet up again to share information. Solid witnesses whether E or H happens, while Dashed witnesses whether F or G happens.

In this example, neither E nor F can happen after the other, and hence Axiom 14 would authorize the formation of the merger $E \cup F$, the event that "either E happens or F happens" at the same time. But this is not an instantaneous event in our sense. The events E and F, although they can both happen, cannot do so simultaneously.

The events E and F in Figure 13 also illustrate the difficulty with Axiom 15. Because neither E nor F can happen after the other, both E^F and F^E are impossible, and hence their merger $E^F \cup F^E$ exists and is equal to the impossible event, Λ. So the axiom is contradicted by the fact that $A \leq E$ and $A \leq F$ hold while $A \leq \Lambda$ does not hold.

When we drop Axioms 14 and 15 in order to accommodate relativity theory, it is no longer guaranteed that two events E and F will have an ending $E \wedge F$ and beginning $E \vee F$, because the mergers in Equations (II.5.1) and (II.5.2) may fail to exist. In Figure 13, for example, neither $E \wedge F$ nor $E \vee F$ exists. The constructions E_5^F and $E_{\rightarrow}^F$ seem to remain reasonable in the relativistic case, and we can still define the E_1^F, E_2^F, E_3^F, and E_5^F, but the events E_3^F and E_4^F can no longer necessarily be interpreted in the way suggested in §6. The construction E^{-}, failure, is not acceptable in the relativistic case, because an event E can be ruled out by events that are incomparable and hence cannot be merged to form a single event.

As we have explained, our complete system of axioms justifies a representation in terms of "point events" that form an event tree; this is the topic of Part V. Without Axioms 14 and 15, we obtain only a partially ordered set of "point events" that is not necessarily a tree, as in Figure 13. Are there other axioms that should be added in the relativistic case? It seems reasonable to add a weaker version of Axiom 14: If div(E,F), the merger $E \cup F$ can be formed. But it is unclear whether other axioms should also be adopted.

In a philosophical study of point events in a relativistic world, Belnap (1992) advances some suggestions for axioms concerning the existence of branching points. Although Belnap's work has provided the impetus for our own thoughts on this topic, we have not adopted his axioms. It is unclear how they can be expressed constructively (rather than as statements about existence) and hence unclear how they can be used to formulate and draw inferences from causal information.

III. Constructive Axioms for Event Spaces

We now provide an axiomatization for event spaces that is constructive in the intuitionistic sense. This means that we avoid the principle of the excluded middle, and we use constructions instead of existential axioms.

We take an intuitionistic approach not because we insist on an intuitionistic philosophy of mathematics, but for more practical reasons:

- The intutionistic emphasis on relations that can actually be observed helps organize the axiomatization in a way that facilitates inferences from observations.
- There are systems, such as Coq and ALF, which translate constructive axiomatizations directly into tools for automated reasoning.
- The discipline provided by the intuitionistic philosophy is an aid to developing concise and effective axiomatizations.

Readers who are unfamiliar or uncomfortable with the constructive viewpoint are encouraged to look first at Part IV, where the axioms we develop in this section are presented in a classical form.

We follow the guidelines for constructive axiomatization developed by Jan von Plato (1995, 1996). We postulate a set of objects (instantaneous events in our case) and predicates and relations for them that may be verified, at least in principle, by finite experience. The relations include a relation of apartness, whose negation is taken as the definition of equality between two objects. We adopt axioms for the predicates and relations that

- reflect the structure of the verification, and
- authorize the substitution of equals for equals in the relations and predicates.

Ideal objects—objects with infinitely precise properties that cannot be verified by direct experience—are introduced by construction. (In geometry, the unique straight line that goes exactly through two given points is the standard example of an ideal object. In our case, an example is the precisely simultaneous happening of two events.) Along with constructed objects, we introduce axioms that

- assert the ideal properties of the constructed objects,
- imply that these ideal properties characterize the constructed objects—i.e., that the constructed objects are the only ones with the properties (for example, the line ln(a,b) constructed from the distinct points a and b is the only line that goes through both a and b), and
- imply that distinct constructions must have distinct inputs, so that equals may be substituted for equals in the construction (for example, ln(a,b) can be distinct from ln(c,d) only if the point a is distinct from the point b or the point c is distinct from the point d).

This differs from the procedure in classical mathematics, where we first assert the existence of an unnamed object with certain properties (we say, for example, that there exists at least one line between any two points), and then we may or may not prove its uniqueness.

In addition to von Plato's guidelines, we adopt an additional methodological principle. The axioms we adopt for each new construction should imply necessary and sufficient conditions, not involving the construction, for whether any relation in the theory holds between the construction and any other object in the theory. This principle gives a certain primacy to the relations in the theory; it suggests that the ideal objects serve only to order and summarize relations among other objects. In some cases, however, it will take us somewhat outside the constructive framework, inasmuch as the necessary and sufficient conditions involve classical existence statements.

The relations that we axiomatize can be thought of as judgments made by Nature, the imagined demigod or super scientist who represents the limits of mortal knowledge. Nature, we suppose, may witness and know everything that any human or artificial scientist might witness or know at a given point in time but lacks God's infinite scope, foresight, and precision of knowledge. When a given event happens, Nature may be aware of its happening and may be aware that certain other events have and have not happened. She may be able to predict at that moment that certain other events will and will not happen in the future. (See Shafer 1996, Chapter 1.) Our constructions define the limits of Nature's potential knowledge and prediction and hence, like the line passing through two given points, do have a meaning of infinite scope and precision. For example, we may construct the event that two events E and F happen simultaneously—an infinitely precise condition on their timing. We may also construct the event that E happens in such a way that F's later happening is possible. When we say that F's happening is possible, we mean that Nature cannot rule it out no matter how far she exploits the immense and perhaps even infinite information available to her.

For brevity, we drop the adjective "instantaneous" when we speak of events, but it will be implicit throughout. An event, as we use the word in this section, happens instantaneously and can happen at most once in the course of events.

The concept of an event space is sufficiently complex that it can be axiomatized in a myriad of ways, no matter whether we take a classical or a constructive approach. Our goal in this section is to produce an intuitionistically acceptable and concise axiomatization based on the two partial orders that we introduced informally in Part II: $\subseteq$ and $\leq$.

A moment's thought reveals that $E \subseteq F$ and $E \leq F$ are not quite appropriate starting points for a constructive axiomatization, for they cannot be verified by finite experience. No amount of experience can tell us that F always happens at exactly the same time whenever E happens, or that F must have happened by the time E happens. Our experience may, on the other hand, authorize the opposite judgements; we may see E happen without F happening at the same time or even without F ever having happened yet. We therefore take the following relations as primitive:

RELATION	Meaning
$E \circ F$	E may happen without F happening simultaneously.
$E \nabla F$	E may happen without F having happened yet.

In §§1-3 we axiomatize the refinement order $\subseteq$. We begin, in §1, by introducing the relation $E \circ F$. We adopt constructive axioms for $\circ$ that make its negation, $\subseteq$, a partial order. In §2 and §3, we adopt two constructions, $E \cap F$ and $E \cup F$, which turn out to be greatest lower and least upper bounds in the partial order $\subseteq$. Our axioms in these three sections are the same as the constructive axioms for lattices formulated by Jan von Plato (1997), except that we authorize the construction of $E \cup F$ only under the assumption that E and F do have at least one upper bound. Because of this restriction, the order $\subseteq$ is not quite a lattice; we obtain a lattice only if we limit our attention to events that refine a particular fixed event.

In §§4-6, we move beyond von Plato's theory in the direction of a Boolean algebra. In §4, we introduce the *resolution axiom*, which implies distributivity and the uniqueness of relative complements. In §5, we introduce the impossible event, which serves as the zero in the lattice. In §6, we introduce relative complements. These additions make our order into a Boolean algebra, except again that not all pairs of events have an upper bound. The axioms for this (almost) Boolean algebra are summarized in §7.

In §§8-15, we turn to our second order, the temporal order $\leq$. We introduce $E \nabla F$ in §8. The denial of $E \nabla F$, $E \leq F$, is the temporal order. We introduce E^F, the part of E that requires F to happen at the same time or earlier, in §9. In §§10-13, we study aspects of the temporal lattice that hold even in a relativistic world. Then, in §§14-15, we impose additional axioms that are appropriate only for a non-relativistic world and lead to the conclusion that the temporal order is a distributive lattice. We summarize the entire axiomatic system in §16.

0. Elements of Intuitionistic Reasoning

The following brief review of the implications of rejecting the principle of the excluded middle should help readers understand how we reason with our axioms. For more systematic and complete expositions of intuitionistic inference, see Dummett (1977), Martin-Löf (1982, 1984), van Dalen (1986, 1997), or Ranta (1994).

The principle of the excluded middle says that for any proposition A, either A or not(A) is true. This principle is adopted in classical logic but not in intuitionistic logic. In intuitionistic logic, asserting a proposition means asserting one has a proof of it, and asserting a disjunction means asserting that one has a proof of one of the disjuncts. We cannot assert "A or not(A)" unless we have a proof of A or of not(A), and we may have neither.

Here are some implications of the intuitionistic rejection of the principle of the excluded middle.

- **not(not(A)) does not imply A.** In the presence of the rules of inference that intuitionists accept, the principle of the excluded middle implies that A and not(not(A)) are equivalent; each implies the other. But once we reject the principle of the excluded middle, this equivalence no longer holds. Intuitionistically, not(A) is the same as the implication $A \Rightarrow \perp$, where $\perp$ is the absurdity. So not(not(A)) is the implication $(A \Rightarrow \perp) \Rightarrow \perp$. This does follow from A; so A implies not(not(A)). But there is no intuitionistic argument taking us the other way, from not(not(A)) to A.
- **Proof by contradiction.** Rejection of the principle of the excluded middle also puts limits on the use of proof by contradiction. Since not(A) means $A \Rightarrow \perp$, we may prove not(A) by assuming A and deriving a contradiction. But if we assume not(A) and derive a contradiction, we have proven only not(not(A)); we have not proven A.
- **Contraposition.** Contraposition is intuitionistically valid; from $A \Rightarrow B$ we may conclude $(B \Rightarrow \perp) \Rightarrow (A \Rightarrow \perp)$, or not(B)$\Rightarrow$not(A). But we cannot conclude $A \Rightarrow B$ from not(B)$\Rightarrow$not(A).
- **Using a disjunction.** What may we conclude from the disjunction A or B? The intuitionistic answer is that we may conclude anything that we may conclude from A and also from B.

- **Eliminating a disjunct.** From the disjunction A or B and the negation not(B), we may derive A intuitionistically. This is because (1) not(B) means $B \Rightarrow \perp$ and (2) $\perp \Rightarrow A$ is accepted as a rule of inference. From A we conclude A. From B and not(B) we conclude $\perp$ and hence A. So from the disjunction A or B and the negation not(B) we conclude A.

1. The Refinement Order: Happening Alone

Sometimes an event E can happen without another event F happening simultaneously. This leaves open whether F has or has not already happened and whether F may or may not happen later. It also leaves open the possibility that E and F may, under different circumstances, happen at the same time.

RELATION	Reading	Example of proof
$E \circ F$	E may happen without F happening simultaneously.	An example where E happens and F does not happen at the same time.

AXIOM	Explanation
1A not($E \circ E$).	E cannot both happen and not happen at the same time.
1B If $E \circ F$, then $E \circ G$ or $G \circ F$.	When E happens without F happening at the same time, either G also happens at that time or else it does not.

Axioms 1A and 1B were proposed by Jan von Plato as general axioms for intuitionistic partial order (von Plato 1997), except that he used a curved inequality sign instead of our $\circ$. He called 1B a *splitting axiom*, and he called any relation satisfying 1A and 1B an *excess* relation. In our case, the excess is only potential; $E \circ F$ means that events *may* turn out in such a way that E exceeds F in the sense that E happens and F does not.

The explanation given for Axiom 1B—that G either does or does not happen—sounds like an appeal to the principle of the excluded middle and therefore requires more elaboration. If Nature can tell that E has happened without F happening at the same time, should she also be able to tell whether G happened at the same time? Our answer is yes, to the exactness required by the axiom. Since Nature's observations are not infinitely precise, her judgment that F did not happen at the same time as E means that there was some finite (as opposed to infinitesimal) interval of time between E's happening and the happening of F, if F happened. If G happens so close to E that Nature, with her merely finite precision, cannot exclude its having happened simultaneously, then it too has happened at a finite distance from any happening of F, and Nature can make the judgment $G \circ F$. Otherwise (if G does not happen or also happens at finite distance from the happening of E) Nature can make the judgment $E \circ G$.[15] See Figure 14.

---

[15] A similar argument is made in the constructive axiomatization of geometry. If two points are distinct in the sense that they are more than infinitesimally far apart, then any third point must be distinct, in the same sense, from at least one of them. See von Plato 1995, p. 173.

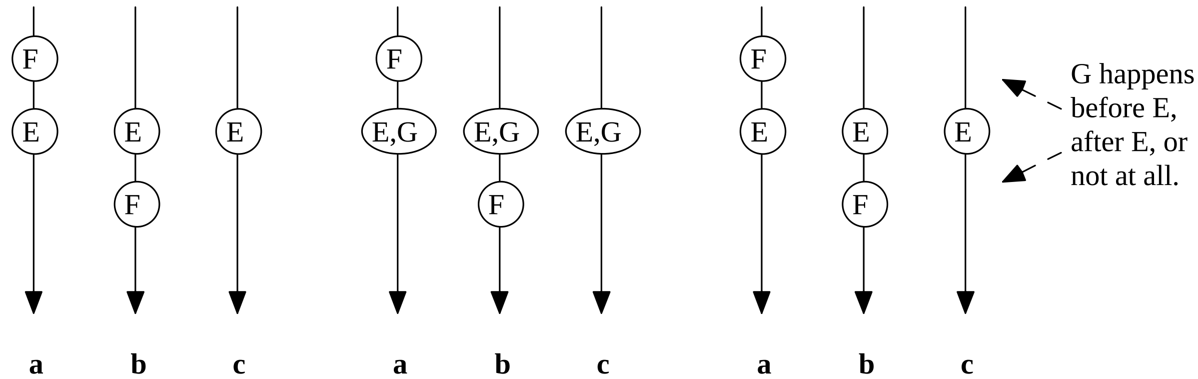

Hypothesis: F, if it happens, happens at a finite distance from E, either before E as in (a) or after E as in (b).

Case 1: G happens infinitesimally close in time to E.

Case 2: G, if it happens, at all, happens at a finite distance from E.

Figure 14 The downward arrows indicate the direction of time. If E $\circ$ F, then one of the three courses of events on the left must be possible: (a) F happens followed by E after a finite (rather than infinitesimal) interval of time, (b) E happens followed by F after a finite interval of time, or (c) E happens and F never happens (at least F never happens while Nature is still watching). In case 1, Nature can make the judgment that G has happened without F happening at the same time, so that G $\circ$ F. In case 2, Nature can make the judgment that E has happened without G happening at the same time, so that E $\circ$ G.

Following von Plato, we use our excess relation to define three other relations: $\neq$, $\subseteq$, and =. The relation $\neq$ is a symmetrization of $\circ$. These relations $\subseteq$ and = are the negations of $\circ$ and $\neq$, respectively. As such, they are less fundamental; they cannot be verified by finite experience. Special cases of E $\subseteq$ F or E = F may be adopted on theoretical grounds (as in 1A), but these relations can be confirmed by experience only in a negative and hence indefinitely protracted way (no example where E happens without F happening at the same time ever turns up).

RELATION	Definition	Reading	Meaning
E $\neq$ F	E $\circ$ F or F $\circ$ E	E is distinct from F.	At least one of the events can happen without the other.
E $\subseteq$ F	not(E $\circ$ F)	E refines F.	Whenever E happens, F happens at the same time.
E = F	E $\subseteq$ F and F $\subseteq$ E	E equals F.	Whenever E or F happens, the other also happens.

We could equivalently define E = F as

- not(E $\neq$ F),
- not(E $\circ$ F or F $\circ$ E), or
- not(E $\circ$ F) and not(F $\circ$ E).

These three assertions are constructively equivalent to each other and to the definition in the table, E $\subseteq$ F and F $\subseteq$ E.

The relations ≠, ⊆, and = have the properties usually associated with the symbols: ≠ is an apartness relation in the sense of Heyting (Dummett 1977, p. 42; van Dalen 1997, p. 179), ⊆ is a partial order, and = is an equivalence relation.

THEOREM	Demonstration
1.1 ≠ is irreflexive: E = E.	Axiom 1A
1.2 ≠ is symmetric: If E ≠ F, then F ≠ E.	Definition of ≠
1.3 If E ≠ F, then E ≠ G or G ≠ F.	Axiom 1B
1.4 ≠ is an apartness relation.	Theorems 1.1, 1.2, and 1.3
1.5 ⊆ is reflexive: E ⊆ E.	Axiom 1A
1.6 ⊆ is transitive: If E ⊆ F and F ⊆ G, then E ⊆ G.	Axiom 1B
1.7 ⊆ is antisymmetric: If E ⊆ F and F ⊆ E, then E = E.	Definition of =
1.8 ⊆ is a partial order.	Theorems 1.5, 1.6, and 1.7
1.9 = is reflexive: E = E.	Theorem 1.1
1.10 = is transitive: If E = F and F = G, then E = G.	Theorem 1.6
1.11 = is symmetric: If E = F, then F = E.	Definition of =
1.12 = is an equivalence relation.	Theorems 1.9, 1.10, and 1.11

Because we are using only intuitionistic principles of inference, it matters that we have defined E = F as not(E ≠ F) rather than defining E ≠ F as not(E = F). The two ways of proceeding are not constructively equivalent. As we have defined the relations, not(E = F) means not(not(E ≠ F)), which is constructively weaker than E ≠ F.

The next two theorems show how the relation E ∘ F respects our partial order.

THEOREM	Demonstration
1.13 If E ∘ F and E ⊆ G, then G ∘ F.	Axiom 1B
1.14 If E ∘ F and G ⊆ F, then E ∘ G.	Axiom 1B

In the classical mathematical treatment of partial order, the meaning of = is implicit. Here we make the meaning of = explicit (two events are equal if neither can happen without the other happening at the same time), and we accept the responsibility of justifying any use we make of =. In addition to verifying that = is an equivalence relation, we must justify any substitution of "equals for equals." The following theorems justify the substitution of equals for equals in the relation E ∘ F.

THEOREM	Demonstration
1.15 If E ∘ F and G = E, then G ∘ F.	Theorem 1.13
1.16 If E ∘ F and G = F, then E ∘ G.	Theorem 1.14

Each time we introduce a predicate, relation, or construction, we will need to verify that we can substitute equals for equals in it. We take it as a principle, in other words, that the axioms accompanying a new predicate, relation, or construction should include axioms that authorize this substitution. This principle applies only to predicates, relations, and constructions that we introduce axiomatically. If we define a new expression (predicate, relation, or construction) in terms of existing expressions, then the validity of the substitution of equals for equals in the new expression will follow from its validity in the existing expressions.

Finally, we note that the equality of two events is determined by the events they can happen without and also by the events that can happen without them.

THEOREM	Demonstration
1.17 If E ∘ G implies F ∘ G, then E ⊆ F.	If we assume E ∘ F, then the hypothesis that E ∘ G implies F ∘ G yields F ∘ F, in contradiction to Axiom 1A.
1.18 If E ∘ G if and only if F ∘ G, then E = F.	Theorem 1.17
1.19 If G ∘ E implies G ∘ F, then F ⊆ E.	If we assume F ∘ E, then the hypothesis that G ∘ E implies G ∘ F yields F ∘ F, in contradiction to Axiom 1A.
1.20 If G ∘ E if and only if G ∘ F, then E = F.	Theorem 1.19

2. Overlap

We now study greatest lower bounds in the partial order ⊆.

RULE OF CONSTRUCTION	Explanation
Overlap From events E and F, construct the event E∩F.	E∩F happens when E and F happen at the same time.

AXIOM	Explanation
2A E∩F ⊆ E and E∩F ⊆ F.	When E∩F happens, E and F both happen.
2B If G ∘ E∩F, then G ∘ E or G ∘ F.	When E and F do not both happen, at least one of them does not happen.

Contraposition of Axiom 2B produces the following more familiar statement.

THEOREM	Demonstration
2.1 If G ⊆ E and G ⊆ F, then G ⊆ E∩F.	Axiom 2B

The overlap E∩F is the greatest lower bound for E and F in the partial order ⊆. (Axiom 2A says it is a lower bound, and Theorem 2.1 says it is greater than or equal to any lower bound.) The operation ∩ therefore has all the algebraic properties of greatest lower bound, some of which we now list.

THEOREM	Demonstration
2.2 ∩ is idempotent: E∩E = E.	An event is its own greatest lower bound.
2.3 ∩ is commutative: E∩F = F∩E.	By the symmetry of the definition of greatest lower bound.
2.4 ∩ is associative: (E∩F)∩G = E∩(F∩G).	Both are the greatest lower bound for the three events E, F, and G.
2.5 If E ⊆ F, then E∩F = E.	E is a lower bound for any event it refines.

We next observe that overlap respects our partial order.

THEOREM	Demonstration
2.6 If E∩F ∘ G∩H, then E ∘ G or F ∘ H.	By Axiom 2B, E∩F ∘ G or E∩F ∘ H. So the conclusion follows from Axiom 2A and Theorem 1.13.
2.7 If E ⊆ G and F ⊆ H, then E∩F ⊆ G∩H.	Theorem 2.6
2.8 If F ⊆ G, then E∩F ⊆ E∩G.	Theorems 2.7 and 1.5
2.9 If E ⊆ G, then E∩F ⊆ G∩F.	Theorems 2.7 and 1.5

Now we consider the principle that equals can be substituted for equals. In the case of constructions, we insist on a constructively stronger principle: the *principle of strong extensionality* (Troelstra and van Dalen 1988, p. 386; von Plato 1997, p. 6). This principle says that if the result of applying a construction to one group of inputs is distinct from the result of applying it to another group of inputs, then the inputs are distinct. Theorem 2.10 says this for the construction E∩F. The three subsequent corollaries, Theorems 2.11, 2.12, and 2.13, express the principle that equals can be substituted for equals in E∩F.

THEOREM	Demonstration
2.10 If E∩F ≠ G∩H, then E ≠ G or F ≠ H.	Theorem 2.6
2.11 If E = G and F = H, then E∩F = G∩H.	Theorem 2.10
2.12 If F = G, then E∩F = E∩G.	Theorems 2.11 and 1.9
2.13 If E = G, then E∩F = G∩F.	Theorems 2.11 and 1.9

We conclude with conditions under which the construction E∩F stands on either side of the relation ∘ with other events.

THEOREM	Demonstration
2.14 G ∘ E∩F if and only if G ∘ E or G ∘ F.	If G ∘ E or G ∘ F, then G ∘ E∩F by Axiom 2A and Theorem 1.14. The opposite implication is Axiom 2B.
2.15 E ∘ F if and only if E ∘ E∩F.	Theorem 2.14 with E for G
2.16 E∩F ∘ G if and only if there exists an event H such that H ⊆ E, H ⊆ F, and H ∘ G.[16]	If E∩F ∘ G, then by Axiom 2A, E∩F is the requisite H. If there is such an H, then H ⊆ E∩F by Theorem 2.1, and hence E∩F ∘ G by Theorem 1.13.

Theorems 2.14 and 2.16 are equivalent, in the presence of Axioms 1A and 1B, to Axioms 2A and 2B.

---

[16] This is a classical existence statement, but the proof reveals the constructive meaning: (1) the three relations H ⊆ E, H ⊆ F, and H ∘ G imply E∩F ∘ G, and (2) if E∩F ∘ G, then E∩F ⊆ E, E∩F ⊆ F, and E∩F ∘ G.

3. Merger

We now study least upper bounds in the partial order $\subseteq$. We continue to follow von Plato's axiomatization for lattices, except that our events do not quite form a lattice, because a least upper bound can be constructed only under the assumption that there is at least one upper bound.

RULE OF CONSTRUCTION	Explanation
Merger From E, F, and I, together with a proof that $E \subseteq I$ and $F \subseteq I$, construct the event $E \cup F$.	$E \cup F$ happens when at least one of the two events E and F happen.

There may be more than one I satisfying $E \subseteq I$ and $F \subseteq I$, but there must be at least one for $E \cup F$ to be formed. Otherwise one of the two events E and F may happen before the other, and in this case they cannot be merged into a single event that may happen at most once.

AXIOM	Explanation
3A If $E \subseteq I$ and $F \subseteq I$, authorizing the construction of $E \cup F$, then $E \subseteq E \cup F$ and $F \subseteq E \cup F$.	When E happens or F happens, $E \cup F$ happens.
3B If $E \cup F \circ G$, then $E \circ G$ or $F \circ G$.	When $E \cup F$ happens, either E happens or F happens.

Contraposition of Axiom 3B produces the following more familiar statement.

THEOREM	Demonstration
3.1 If $E \subseteq G$ and $F \subseteq G$, then $E \cup F \subseteq G$.	Axiom 3B

The merger $E \cup F$ is the least upper bound for E and F in the partial order $\subseteq$. Axiom 3A says it is an upper bound, and Theorem 3.1 says it is less than or equal to any upper bound. The operation $\cup$ has all the algebraic properties of least upper bound, some of which we now list.

THEOREM	Demonstration
3.2 $\cup$ is idempotent: $E \cup E = E$.	An event is its own least upper bound.
3.3 $\cup$ is commutative: If $E \subseteq I$ and $F \subseteq I$, authorizing the construction of $E \cup F$ and $F \cup E$, then $E \cup F = F \cup E$.	By the symmetry of the definition of least upper bound.
3.4 $\cup$ is associative: If $E \subseteq I$, $F \subseteq I$, and $G \subseteq I$, authorizing the construction of $(E \cup F) \cup G$ and $E \cup (F \cup G)$, then $(E \cup F) \cup G = E \cup (F \cup G)$.	Both are the least upper bound for the three events E, F, and G.
3.5 If $E \subseteq F$, then $E \cup F = F$.	An event is an upper bound for any refinement.

Merger respects our partial order. (We leave implicit the assumptions required to authorize the construction of the mergers in these theorems.)

THEOREM	Demonstration
3.6 If $E \cup F \circ G \cup H$, then $E \circ G$ or $F \circ H$.	By Axiom 3B, $E \circ G \cup H$ or $F \circ G \cup H$. So the conclusion follows from Axiom 3A and Theorem 1.14.
3.7 If $E \subseteq G$ and $F \subseteq H$, then $E \cup F \subseteq G \cup H$.	Theorem 3.6
3.8 If $F \subseteq G$, then $E \cup F \subseteq E \cup G$.	Theorems 3.7 and 1.5
3.9 If $E \subseteq G$, then $E \cup F \subseteq G \cup F$.	Theorems 3.7 and 1.5

Merger obeys strong extensionality, so that we can substitute equals for equals in $E \cup F$. (We again leave implicit the assumptions required for the construction of the mergers.)

THEOREM	Demonstration
3.10 If $E \cup F \neq G \cup H$, then $E \neq G$ or $F \neq H$.	Theorem 3.6
3.11 If $E = G$ and $F = H$, then $E \cup F = G \cup H$.	Theorem 3.10
3.12 If $F = G$, then $E \cup F = E \cup G$.	Theorems 3.11 and 1.9
3.13 If $E = G$, then $E \cup F = G \cup F$.	Theorems 3.11 and 1.9

A partially ordered set in which each pair of elements has a least upper bound and a greatest lower bound is called a *lattice*. Our events do not form a lattice in their totality under the partial ordering $\subseteq$, because upper bounds to do not exist for all pairs of events, but we do have a lattice when we limit our attention to events that refine a particular event (say all E such that $E \subseteq I$).

Lattices can also be characterized algebraically; instead of beginning with the idea that $\cap$ and $\cup$ represent greatest lower bound and least upper bound, respectively, one postulates that they obey certain algebraic axioms (Davey and Priestley 1990, Chapter 5). Of all the standard algebraic axioms, the absorption laws are the only ones we have not yet derived. We now derive them.

THEOREM	Demonstration
3.14 $E \cup (E \cap F) = E$.	E is the least upper bound for $E \cap F$ and E.
3.15 If $E \subseteq I$ and $F \subseteq I$, authorizing the construction of $E \cup F$, then $E \cap (E \cup F) = E$.	E is the greatest lower bound for $E \cup F$ and E.

We conclude with conditions under which the construction $E \cup F$ stands on either side of the relation $\circ$ with other events.

Suppose $E \cup F$ can be constructed.	
THEOREM	Demonstration
3.16 $E \cup F \circ G$ if and only if $E \circ G$ or $F \circ G$.	If $E \circ G$ or $F \circ G$, then $E \cup F \circ G$ by Axiom 3A and Theorem 1.13. The opposite implication is Axiom 3B.
3.17 $E \circ F$ if and only if $E \cup F \circ F$.	Theorem 3.16 with F for G

3.18 G ∘ E∪F if and only if there exists an event H such that E ⊆ H, F ⊆ H, and G ∘ H.	If G ∘ E∪F, then by Axiom 3A, E∪F is the requisite H. If there is such an H, then E∪F ⊆ H by Theorem 3.1, and hence G ∘ E∪F by Theorem 1.14.

Theorems 3.16 and 3.18 are equivalent, in the presence of Axioms 1A and 1B, to Axioms 3A and 3B.

4. The Resolution Axiom

In this section, we adopt an axiom, the *resolution axiom*, that assures that the refinement order is distributive and that relative complements in it are unique.

One way to motivate the resolution axiom is to argue for a strengthening of Axiom 1B, the splitting axiom. The splitting axiom says that if E ∘ F, then E ∘ G or G ∘ F. The reason is that either G does not happen at the same time as E (so that E happens without either F or G happening) or else it does (so that E and G both happen without F happening). Now that we have the concepts of overlap and merger to work with, we might express this more strongly:

If E ∘ F, then E ∘ F∪G or E∩G ∘ F.

But this may be too strong for our intuitionistic concept of Nature. Nature may be unable to tell whether happenings of E and G are precisely simultaneous. We also have the problem that the hypothesis E ∘ F does not imply the existence of F∪G.

By putting the proposed axiom in contrapositive form, however, we obtain an intuitionistically acceptable elaboration of the meaning of the constructions F∪G and E∩G.

AXIOM	Explanation
4 If E ⊆ F∪G and E∩G ⊆ F, then E ⊆ F.	Suppose E happens. If this implies that F or G happens at the same time, and also that G's happening at the same time implies F happens at the same time, then it implies F happens at the same time.

(Here the existence of F∪G is implicit in the hypothesis.)

Axiom 4 is a powerful tool for proving inequalities; for it allows us to deduce E ⊆ F from two related but weaker inequalities. If we shift the conversation from events to propositions (interpreting ⊆, ∪, and ∩ as implication, disjunction, and conjunction, respectively), then this axiom expresses a simple and familiar tactic in mathematical reasoning: when we are trying to prove F and find we can prove only F or G, we adopt G as an additional assumption and try again to prove F. We call the axiom the resolution axiom because in this context it is a constructive version of the method of resolution for theorem proving: from the clauses {not(E),F,G} and {not(E),F,not(G)}, infer the clause {not(E),F} (Robinson 1965).

As we will see, the resolution axiom implies modularity (Theorem 4.4), distributivity (Theorems 4.5 and 4.6), and the uniqueness of relative complements when they exist (Theorem 6.1).

A lattice is said to be *distributive* if it satisfies the two distributive laws:

1. A∪(B∩C) = (A∪B)∩(A∪C).
2. A∩(B∪C) = (A∩B)∪(A∩C).

It is said to be *modular* if it satisfies a weaker condition called the modular law:

3. If A ⊆ C, then A∪(B∩C) = (A∪B)∩C.

The resolution axiom guarantees that these laws hold whenever the mergers in them can be constructed.

For clarity, let us begin by noting that half of each law can be derived without using the resolution axiom.

Suppose A∪B can be constructed.	
THEOREM	Demonstration
4.1 If A ⊆ C, then A∪(B∩C) ⊆ (A∪B)∩C.	Both A and B∩C refine both A∪B and C. So the inequality follows from the definitions of least upper and greatest lower bound.

Suppose A∪B and A∪C can be constructed.	
THEOREM	Demonstration
4.2 A∪(B∩C) ⊆ (A∪B)∩(A∪C).	Both A and B∩C refine both A∪B and A∪C. So the inequality follows from the definitions of least upper bound and greatest lower bound.

Suppose B∪C can be constructed.	
THEOREM	Demonstration
4.3 (A∩B)∪(A∩C) ⊆ A∩(B∪C).	Both A∩B and A∩C refine both A and B∪C. So the inequality follows from the definitions of least upper bound and greatest lower bound.

Now we use the resolution axiom to prove the remaining half of each of law. We begin with the modular law (without the assumption A ⊆ C, which is not needed for this half).

Suppose A∪B can be constructed.	
THEOREM	Demonstration
4.4 (A∪B)∩C ⊆ A∪(B∩C).	Substitute (A∪B)∩C for E, A∪(B∩C) for F, and B for G in Axiom 4.

Now we prove the remaining half of the first distributive law.

Suppose A∪B and A∪C can be constructed.	
THEOREM	Demonstration
4.5 (A∪B)∩(A∪C) ⊆ A∪(B∩C).	Substitute (A∪B)∩(A∪C) for E, A∪(B∩C) for F, and B for G in Axiom 4. This reduces the problem to an instance of Theorem 4.4.

Finally, we prove the remaining half of the second distributive law.

Suppose B∪C can be constructed.	
THEOREM	Demonstration
4.6 A∩(B∪C) ⊆ (A∩B)∪(A∩C).	Substitute A∩(B∪C) for E, (A∩B)∪(A∩C) for F, and B for G in Axiom 4. Again, this reduces the problem to an instance of Theorem 4.4.

We may now say that the lattice formed by the refinements of a fixed event is distributive.

5. Possibility and Impossibility

An element E_0 in a partial order is a *zero* if $E_0 \subseteq E$ for all E. It follows from this definition that any two zeroes for a given partial order are equal, and so we speak of *the* zero when it exists.

We now construct the impossible event.

RULE OF CONSTRUCTION	Explanation
Impossible Event Construct the event Λ.	Λ is the impossible event.

We adopt a single axiom for Λ, which tells us that it is the zero for the refinement order.

AXIOM	Explanation
5 Λ ⊆ E.	The impossible event cannot happen.

THEOREM	Demonstration
5.1 If E ◦ F, then E ◦ Λ.	Axiom 5 and Theorem 1.14
5.2 E∩Λ = Λ.	Axiom 5 and Theorem 2.5
5.3 E∪Λ = E.	Axiom 5 and Theorem 3.5

Axiom 5 suffices to specify conditions under which the construction Λ stands on either side of the relation ◦ with other events. The relation Λ ◦ E never holds (Axiom 5), and the relation E ◦ Λ holds if and only if E ◦ F for some F (Theorem 5.1).

Because Λ ⊆ E for all E, the condition E ⊆ Λ is equivalent to E = Λ.

We define possibility and impossibility in terms of the impossible event.

PREDICATE	Definition	Reading	Meaning
poss(E)	E ◦ Λ	E is possible.	E may happen.
imposs(E)	not(poss(E))	E is impossible.	E cannot happen.

We call a possible event a *situation*.

Here are a few facts about possibility.

THEOREM	Demonstration
5.4 imposs(E) if and only if E = Λ.	Axiom 5 and the definition of imposs(E)

5.5 poss(E) if and only if $E \circ F$ for some F.	Theorem 5.1 and the definition of poss(E)
5.6 If poss(E) and $E \subseteq F$, then poss(F).	Theorem 1.13 and the definition of poss(E)

We define two binary predicates, lap(E,F) and dis(E,F).

RELATION	Definition	Reading	Meaning
lap(E,F)	$E \cap F \circ \Lambda$	E and F overlap.	E and F may happen at the same time.
dis(E,F)	$E \cap F = \Lambda$	E and F are disjoint.	E and F cannot happen at the same time.

The relations lap(E,F) and dis(E,F) are symmetric; lap(F,E) if and only if lap(E,F), and dis(F,E) if and only if dis(E,F). We will use this symmetry without comment.

Here are some implications of disjointness.

THEOREM	Demonstration
5.7 If dis(E,F) and poss(G), then $G \circ E$ or $G \circ F$.	The assumptions $E \cap F = \Lambda$ and $G \circ \Lambda$ imply $G \circ E \cap F$. Axiom 2B then implies that $G \circ E$ or $G \circ F$.
5.8 If dis(E,F) and $G \subseteq F$, then dis(E,G).	By Theorem 2.8, $G \subseteq F$ implies $E \cap G \subseteq E \cap F$. So $E \cap F = \Lambda$ implies $E \cap G = \Lambda$.

6. Relative Complements in the Refinement Order

We now study relative complements in the partial order $\subseteq$. The axioms we adopt will assure that the refinements of a possible event form a Boolean algebra.

We say that G is a *complement* of F relative to E if

$$(E \cap F) \cap G = \Lambda \text{ and } (E \cap F) \cup G = E.$$

If F is a refinement of E, then this condition simplifies to

$$F \cap G = \Lambda \text{ and } F \cup G = E.$$

The general concept can be understood in terms of the special case: G is a complement of F relative to E if and only if it is a complement of $E \cap F$ relative to E.

The following theorem tells us that relative complements are unique when they exist.

THEOREM	Demonstration
6.1 If $A \cap C = A \cap D$ and $A \cup C = A \cup D$, then C = D.	Substituting C for E, D for F, and A for G in Axiom 4, we obtain $C \subseteq D$. By symmetry, $D \subseteq C$.

So we may call E\F *the* complement of F relative to E.

We now authorize a new construction and adopt axioms that guarantee it is the relative complement.

RULE OF CONSTRUCTION	Explanation
Relative Complement From E and F, construct E\F.	E\F is the event that E happens without F happening at the same time.

AXIOM	Explanation
6A E\F ∘ G if and only if E ∘ (E∩F)∪(E∩G).	Both propositions assert that E can happen with neither F nor G happening at the same time.
6B G ∘ E\F if and only if G ∘ E or E∩F∩G ∘ Λ.	There are two ways G can happen without E\F happening at the same time. One is for E not to happen then (this proves G ∘ E). The other is for E and F both to happen then (this proves E∩F∩G ∘ Λ).

These axioms state directly conditions under which the construction E\F stands on either side of the relation ∘ with other events.

We now establish that E\F is indeed the complement of E relative to F—the unique refinement of E that is disjoint from E∩F and has E as its merger with E∩F.

THEOREM	Demonstration
6.2 E\F ⊆ E.	Axiom 6B, with E\F substituted for G, together with Axiom 1A
6.3 (E∩F)∩(E\F) = Λ.	
6.4 (E∩F)∪(E\F) = E.	The relation (E∩F)∪(E\F) ⊆ E follows from Axiom 2A and Theorems 6.2 and 3.1. To establish E ⊆ (E∩F)∪(E\F), we assume E ∘ (E∩F)∪(E\F), rewrite this as E ∘ (E∩F)∪(E∩(E\F)) using Theorems 6.2 and 2.5, and then deduce from Axiom 6A, with E\F for G, that E\F ∘ E\F, in contradiction to Axiom 1A.

As we noted before proving Theorem 6.1, it follows from the definition of relative complement that the complement of E relative to F is the same as its complement relative to E∩F. This is recorded by the following theorem.

THEOREM
6.5 E\F = E\(E∩F).

Here are some additional properties of relative complements.

THEOREM	Demonstration
6.6 E\F ∘ G if and only if E\G ∘ F.	Axiom 6A and Theorem 3.3
6.7 E\F ⊆ G if and only if E\G ⊆ F	Theorem 6.6
6.8 Suppose F∪G can be constructed. Then E\F ∘ G if and only if E ∘ F∪G.	When F∪G can be constructed, (E∩F)∪(E∩G) is equal to E∩(F∪G) by distributivity. So E\F ∘ G is equivalent to E ∘ E∩(F∪G) by Axiom 6A. And E ∘ E∩(F∪G) is equivalent to E ∘ F∪G by Theorem 2.15.
6.9 E\E = Λ.	Axiom 6A with E for F and Λ for G
6.10 E\Λ = E.	Theorem 6.4 with Λ for F
6.11 (E\F)∩F = Λ.	Since E\F ⊆ E (Theorem 6.2), (E\F)∩F = (E\F)∩F∩E (Theorem 2.5), which is equal to Λ by Theorem 6.3.
6.12 If E\F ∘ G, then E ∘ G.	Theorems 6.2 and 1.13
6.13 If (E\F)∩G ∘ Λ, then G ∘ F.	By Theorem 6.11, (E\F)∩G ∘ Λ implies (E\F)∩G ∘ (E\F)∩F, and by Theorem 2.6, this implies G ∘ F.

6.14 E\F ∘ Λ if and only if E ∘ F.	Axiom 6A and Theorem 2.15
6.15 E\F = Λ if and only if E ⊆ F.	Theorem 6.14

The construction E\F is monotonic with respect to the partial order ⊆ in both its arguments, in opposite directions.

THEOREM	Demonstration
6.16 If G\H ∘ E\F, then G ∘ E or F ∘ H.	Suppose G\H ∘ E\F. By Axiom 6B, G\H ∘ E (which implies G ∘ E by Theorem 6.12) or E∩F∩(G\H) ∘ Λ (which implies F ∘ H by Theorem 6.13).
6.17 If E ⊆ G, then E\F ⊆ G\F.	Theorem 6.16
6.18 If G ⊆ F, then E\F ⊆ E\G.	Theorem 6.16

We have strong extensionality for E\F, and so we can substitute equals for equals.

THEOREM	Demonstration
6.19 If E\F ≠ G\H, then E ≠ G or F ≠ H.	Theorem 6.16
6.20 If E = G, then E\F = G\F.	Theorem 6.17
6.21 If G = F, then E\F = E\G.	Theorem 6.18

Our definition of relative complement—E\F is the unique event that is disjoint from F and has E as its merger with E∩F—is algebraic. The event E\F can also be characterized, however, in a more order-theoretic way. It is the largest refinement of E that is disjoint from F. This is established as follows.

THEOREM	Demonstration
6.22 If G ⊆ E and F∩G = Λ, then G ⊆ E\F.	Axiom 6B
6.23 E\F is the largest refinement of E that is disjoint from F.	This is the content of Theorems 6.2, 6.11, and 6.22.

Moreover, (E\F)\G is the largest refinement of E that is disjoint from both F and G. To see this, we reason as follows.

THEOREM	Demonstration
6.24 (E\F)\G ⊆ E.	Theorem 6.2
6.25 ((E\F)\G)∩G = Λ.	Theorem 6.11
6.26 ((E\F)\G)∩F = Λ.	Since E\F is disjoint from F (Theorem 6.11), and (E\F)\G ⊆ E\F (Theorem 6.2), (E\F)\G is disjoint from F (Theorem 5.8).
6.27 If H ⊆ E, H∩F = Λ, and H∩G = Λ, then H ⊆ (E\F)\G.	By Theorem 6.22, it follows from H⊆ E and H∩F = Λ that H ⊆ E\F. And it follows from H ⊆ E\F and H∩G = Λ that H ⊆ (E\F)\G.
6.28 (E\F)\G is the largest refinement of E that is disjoint from both F and G.	This is the content of Theorems 6.24, 6.25, 6.26, and 6.27.

These theorems follow.

THEOREM	Demonstration
6.29 (E\F)\F = E\F.	Theorems 6.23 and 6.28
6.30 (E\F)\G = (E\G)\F.	Theorem 6.28

By Theorems 6.15 and 6.30, the four following statements are all equivalent: E\F ⊆ G, E\G ⊆ F, (E\F)\G = Λ, and (E\G)\F = Λ.

Here are some further results that will prove helpful.

THEOREM	Demonstration
6.31 If G ⊆ E, then G∩(E\F) = G\F.	By Theorems 6.2, 6.17, and 2.1, G\F ⊆ G∩(E\F). By Axiom 2A and Theorems 6.11 and 6.23, G∩(E\F) ⊆ G\F.
6.32 If E\F ∘ (E\A)\B, then A ∘ F or B ∘ F.	Axiom 6B and Theorems 6.13 and 6.16
6.33 If A ⊆ F and B ⊆ F, then E\F ⊆ (E\A)\B.	Theorem 6.32

Here is a theorem that does not involve relative complements in its statement but which apparently requires the axioms for relative complements for its proof.

THEOREM	Demonstration
6.34 If E ∘ Λ, then lap(E,F) or E ∘ F.	Theorem 6.4 says that (E∩F)∪(E\F) = E. If E ∘ Λ, then Axiom 3B yields the conclusion that E∩F ∘ Λ (i.e., lap(E,F)) or E\F ∘ Λ (whence E ∘ F by Theorem 6.14).

This leads to the following symmetric interpretation of (H\E)\F being impossible.

THEOREM	Demonstration
6.35 (H\E)\F = Λ if and only if [poss(G) and G ⊆ H] implies [lap(G,E) or lap(G,F)].	Suppose [poss(G) and G ⊆ H] does imply [G∩E ∘ Λ or G∩F ∘ Λ]. Taking (H\E)\F for G, we see that (H\E)\F ∘ Λ would imply ((H\E)\F)∩E ∘ Λ or ((H\E)\F)∩F ∘ Λ, in contradiction to Theorems 6.25 and 6.26. Going the other way, suppose (H\E)\F = Λ, poss(G), and G ⊆ H. By Theorem 6.15, H\E ⊆ F. By Theorem 6.34, G ∘ H\E (which implies G∩E ∘ Λ by Axiom 6B) or G∩(H\E) ∘ Λ (which implies G∩F ∘ Λ).

Intuitively, (H\E)\F = Λ means that H's happening must involve either E's happening or F's happening. Theorem 6.35 yields this interesting implication: If G ⊆ H and G are disjoint from both E and F, then G is impossible.

We now show that our other constructions, E∩F, E∪F, and Λ, can be defined in terms of relative complement. First, E∩F = E\(E\F).

THEOREM	Demonstration
6.36 E∩F = E\(E\F).	If E\(E\F) ∘ E∩F, then Axiom 6A, gives E ∘ (E\F)∪(E∩F), contradicting Theorem 6.4. If E∩F ∘ E\(E\F), Axiom 6B gives either E∩F ∘ E, contradicting Axiom 2A, or (E∩F)∩(E\F) ∘ Λ, contradicting Theorem 6.3.

Second, E∪F = I\((I\E)\F), where I is any event refined by both E and F.

Suppose E ⊆ I and F ⊆ I.	
THEOREM	Demonstration
6.37 F ⊆ I\((I\E)\F).	Axiom 6B and Theorem 6.11
6.38 E ⊆ I\((I\E)\F).	Theorem 6.30 and 6.37
6.39 If E ⊆ G, and F ⊆ G, then I\((I\E)\F) ⊆ G.	By Theorem 6.33, I\G ⊆ (I\E)\F. So by Theorem 6.7, I\((I\E)\F) ⊆ G.
6.40 E∪F = I\((I\E)\F).	By Theorems 6.37, 6.38, and 6.39, I\((I\E)\F) is the least upper bound of E and F in the refinement order.

Third, the impossible event is equal to E\E for any event E.

THEOREM	Demonstration
6.41 Λ = E\E.	Theorem 6.9

Finally, we prove DeMorgan's laws.

Suppose E ⊆ I and F ⊆ I.
THEOREM
6.42 I\(E∩F) = (I\E)∪(I\F).
6.43 I\(E∪F) = (I\E)∩(I\F).

To prove Theorem 6.42, we first use E∩F ⊆ E and E∩F ⊆ F to conclude, by Theorem 6.18, that I\E ⊆ I\(E∩F) and I\F ⊆ I\(E∩F), so that (I\E)∪(I\F) ⊆ I\(E∩F). To prove the opposite refinement, we use Theorem 6.17 to obtain E\F ⊆ I\F and then Theorem 6.18 to obtain E\(I\F) ⊆ E\(E\F), whence, by Theorem 6.36, E\(I\F) ⊆ E∩F. Because E = I\(I\E) (by Theorem 6.36), we obtain (I\(I\E))\(I\F) ⊆ E∩F. Another application of Theorem 6.18 yields I\(E∩F) ⊆ I\((I\(I\E))\(I\F)). By Theorem 6.40, this can be written I\(E∩F) ⊆ (I\E)∪(I\F).

To prove Theorem 6.43, we first use E ⊆ E∪F and F ⊆ E∪F to conclude, by Theorem 6.18, that I\(E∪F) ⊆ I\E and I\(E∪F) ⊆ I\F, so that I\(E∪F) ⊆ (I\E)∩(I\F). To show the opposite refinement, we use (I\E)∩(I\F) ⊆ (I\E)\F to obtain I\((I\E)\F) ⊆ I\((I\E)∩(I\F)), or E∪F ⊆ I\((I\E)∩(I\F)), which yields (I\E)∩(I\F) ⊆ I\(E∪F) by Theorem 6.7.

7. The Refinement Order: Summary

Here is a summary of the axioms and primitive relations and constructions we have introduced so far.

Relation
E ∘ F

We write E ⊆ F for not(E ∘ F).

RULE OF CONSTRUCTION
Overlap From events E and F, construct the event E∩F.
Merger From events E, F, and I and proofs of E ⊆ I and F ⊆ I, construct the event E∪F.
Impossible Event Construct the event Λ.
Relative Complement From events E and F, construct the event E\F.

AXIOM
1A E ⊆ E.
1B If E ∘ F, then E ∘ G or G ∘ F.
2A E∩F ⊆ E and E∩F ⊆ F.
2B If G ∘ E∩F, then G ∘ E or G ∘ F.
3A If E∪F can be constructed, then E ⊆ E∪F and F ⊆ E∪F.
3B If E∪F ∘ G, then E ∘ G or F ∘ G.
4 If E ⊆ F∪G and E∩G ⊆ F, then E ⊆ F.
5 Λ⊆ E.
6A E\F ∘ G if and only if E ∘ (E∩F)∪(E∩G).
6B G ∘ E\F if and only if G ∘ E or E∩F∩G ∘ Λ.

We could make this system more parsimonious, if we wished, by defining the other constructions in terms of the relative complement, as in Theorems 6.36, 6.40, and 6.41. However, this would make the axioms much less readable.

Recall that a *Boolean algebra* is a complemented distributive lattice with a zero (an element Λ such that Λ ⊆ E for all E) and a unit (an element Ω such that E ⊆ Ω for all E). Our axioms postulate the existence of a zero but do not postulate the existence of a unit. They also do not quite give a lattice, since two events may fail to have a merger. If we were to add the construction Ω and the axiom E ⊆ Ω to our system, a merger could be constructed for any pair of events, and so we would have a constructive system of axioms for the concept of a Boolean algebra.[17] In any case, our axioms as they stand imply that the refinements of a fixed possible event E form a Boolean algebra, with E as its unit.

8. The Temporal Order: Happening First

We now turn to the temporal ordering of events—the possibility that one may occur after another. We begin with another excess relation.

[17] As far as we are aware, no constructive system of axioms for Boolean algebras has previously been published. We are also not aware of any previous use of our resolution axiom, Axiom 4, in any axiomatization for Boolean algebras, constructive or classical.

RELATION	Reading	Example of proof
E ∇ F	E may happen without F ever having happened.	An example where E has happened and F has not.

AXIOM	Explanation
8A If E ∇ F, then E ∘ F.	If E has happened and F has not, then F did not happen at the same time E happened.
8B If E ∇ F, then E ∇ G or G ∇ F.	Suppose E happens without F yet having happened. Then G has either happened or not. In the first case, G has happened and F has not. In the second case, E has happened and G has not.
8C If E ∘ F, then E ∇ F or F ∇ G or E ∘ G.	Suppose E happens without F happening at the same time. If G did not happen as E happened, E ∘ G. If F did not happen earlier, E ∇ F. If G did happen at the same time and F happened earlier, then F happened without G ever having happened, for otherwise G would happen twice.
8D If E ⊆ I and F ⊆ I, then G ∇ E∩F implies G ∇ E or G ∇ F.	Because E and F are refinements of the same instantaneous event, they can both happen only by happening at the same time.

Axiom 8D requires the relation ∇ to respect the instantaneous nature of the event I.

The following figures lay out the constructive intuition behind the specific instances of the principle of the excluded middle embodied in Axioms 8B and 8C.

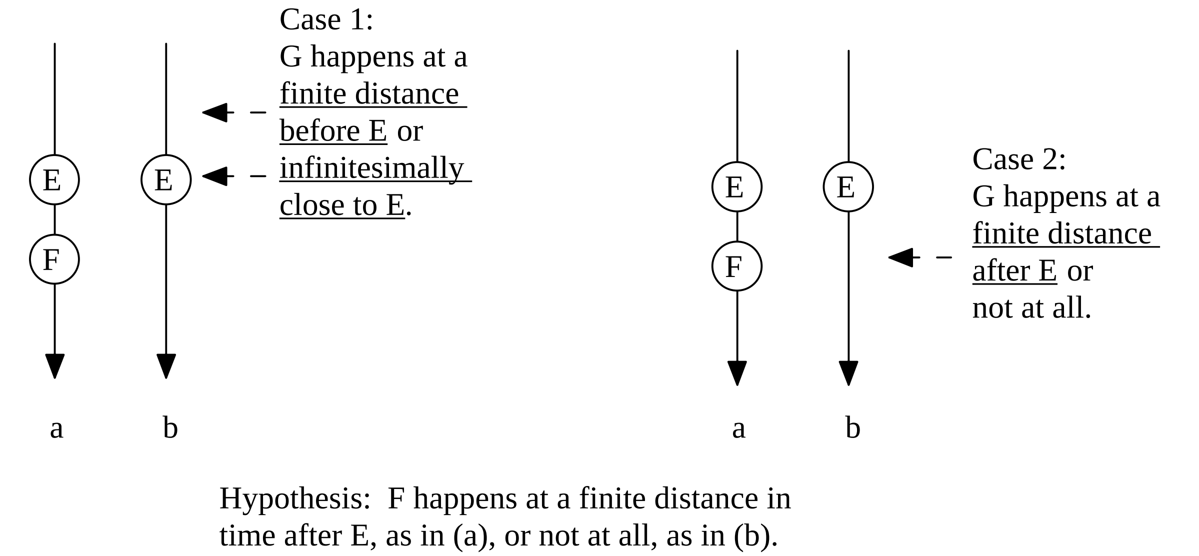

Figure 15 An explanation of Axiom 8B. If E ∇ F, then one of two courses of events must be possible: (a) E happens followed by F after a finite interval of time, or (b) E happens and F never happens. In case 1, Nature can make the judgment that G has happened without F yet having happened, so that G ∇ F. In case 2, Nature can make the judgment that E has happened without G yet having happened, so that E ∇ G.

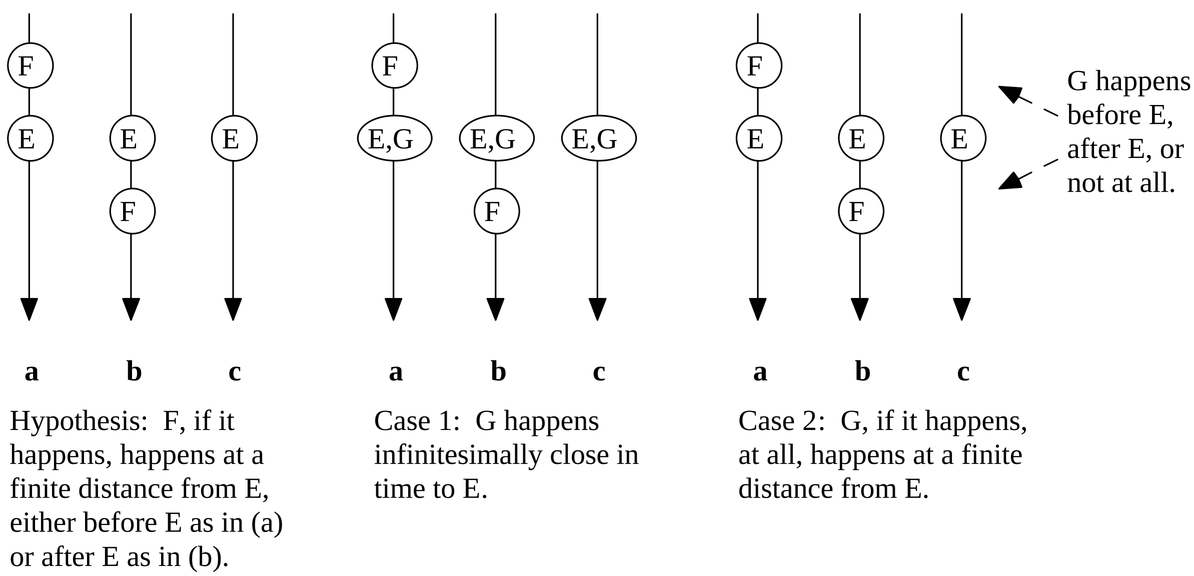

Figure 16 To explain Axiom 8C, we use the same picture that we used to explain Axiom 1B. If E ∘ F, then one of the three courses of events on the left must be possible: (a) F happens followed by E after a finite interval of time, (b) E happens followed by F after a finite interval of time, or (c) E happens and F never happens. In case 1a, Nature can make the judgment that F has happened without G yet having happened, so that F ∇ G. In cases 1b and 1c, Nature can make the judgment that E has happened without F yet having happened, so that E ∇ F. In case 2, Nature can make the judgment that E has happened without G happening at the same time, so that E ∘ G.

THEOREM	Demonstration
8.1 not(E ∇ E).	Axioms 8A and 1A
8.2 If E ∘ F, then E ∇ F or F ∇ E.	Axiom 8C with E for G and Axiom 1A
8.3 If E ∘ Λ, then E ∇ Λ.	By Theorem 8.2, E ∘ Λ implies E ∇ Λ or Λ ∇ E. But by Axiom 8A, Λ ∇ E implies Λ ∘ E, contradicting Axiom 5.
8.4 [E ∘ F or F ∘ E] if and only if [E ∇ F or F ∇ E].	Axiom 8A and Theorem 8.2

Theorem 8.1 and Axiom 8B tell us that ∇, like ∘, is an excess relation. So we can draw all the conclusions about ∇ that we drew about ∘ in §1.

Theorem 8.4 tells us that the excess relation ∇ produces the same apartness relation ≠ and hence the same equality relation = as the excess relation ∘ produces. (Recall from §1 that E ≠ F means, by definition, that E ∘ F or F ∘ E, while E = F means not(E ≠ F).)

We now introduce the relation that denies E ∇ F.

RELATION	Definition	Reading	Meaning
E ≤ F	not(E ∇ F)	E requires F.	E can happen only if F happens at the same time or earlier.

Because it is the negation of an excess relation, ≤ is a partial order. We call it the *temporal order*. As we note in Theorem 8.5, E ≤ F holds whenever E ⊆ F does. As we note in Theorem 8.7, ≤ coincides with ⊆ when we consider only refinements of a fixed event. (This is why ≤ points in the direction it does, which might otherwise seem odd.)

THEOREM	Demonstration
8.5 If E ⊆ F, then E ≤ F.	Axiom 8A
8.6 If E ⊆ I and F ⊆ I, then E ∘ F implies E ∇ F.	Axiom 8C with I for G
8.7 If E ⊆ I and F ⊆ I, then E ≤ F implies E ⊆ F.	Theorem 8.6

It follows from Theorem 8.5 that Λ, the zero for ⊆, is also the zero for ≤.

The next two theorems are exactly analogous to Theorems 1.13 and 1.14; they show that ∇ respects the partial order ≤ just as ∘ respects the partial order ⊆.

THEOREM	Demonstration
8.8 If E ∇ F and E ≤ G, then G ∇ F.	Axiom 8B
8.9 If E ∇ F and G ≤ F, then E ∇ G.	Axiom 8B

Hence we can substitute equals for equals in E ∇ F.

THEOREM	Demonstration
8.10 If E ∇ F and G = E, then G ∇ F.	Theorem 8.8
8.11 If E ∇ F and G = F, then E ∇ G.	Theorem 8.9

Next, we note that ∇ respects ⊆ in the same way as it respects ≤.

THEOREM	Demonstration
8.12 If E ∇ F and E $\subseteq$ G, then G ∇ F.	Theorems 8.5 and 8.8
8.13 If E ∇ F and G $\subseteq$ F, then E ∇ G.	Theorems 8.5 and 8.9

Contraposing these last two theorems, we obtain the following statements, which can also be thought of as aspects of the transitivity of $\leq$.

THEOREM	Demonstration
8.14 If E $\subseteq$ G and G $\leq$ F, then E $\leq$ F	Theorem 8.12
8.15 If E $\leq$ G and G $\subseteq$ F, then E $\leq$ F.	Theorem 8.13

When do the constructions we have already introduced, E$\cap$F, E$\cup$F, Λ, and E\F, stand on one or the other side of the relation ∇ with another event? One way to answer this question is to use the characterizations of these constructions in terms of the order $\subseteq$. This is straightforward, but we spell it out for the sake of completeness.

THEOREM	Demonstration
8.16 E$\cap$F ∇ G if and only if there exists an event H such that H $\subseteq$ E, H $\subseteq$ F, and H ∇ G.	If E$\cap$F ∇ G, then E$\cap$F is the requisite H. Going the other way, suppose H $\subseteq$ E and H $\subseteq$ F. Then H $\subseteq$ E$\cap$F by Theorem 2.1. This, together with H ∇ G, implies E$\cap$F ∇ G by Theorem 8.12.
8.17 G ∇ E$\cap$F if and only if [H $\subseteq$ E and H $\subseteq$ F imply G ∇ H].	Suppose G ∇ E$\cap$F. From H $\subseteq$ E and H $\subseteq$ F, we obtain H $\subseteq$ E$\cap$F by Theorem 2.1 and then G ∇ H by Theorem 8.13. Going the other way, suppose H $\subseteq$ E and H $\subseteq$ F imply G ∇ H. By Axiom 2A, E$\cap$F $\subseteq$ E and E$\cap$F $\subseteq$ F. So G ∇ E$\cap$F.
8.18 Suppose E$\cup$F can be constructed. Then E$\cup$F ∇ G if and only if [E $\subseteq$ H and F $\subseteq$ H imply H ∇ G].	Suppose E$\cup$F ∇ G. From E $\subseteq$ H and F $\subseteq$ H, we obtain E$\cup$F $\subseteq$ H by Theorem 3.1 and then H ∇ G by Theorem 8.12. Going the other way, suppose E $\subseteq$ H and F $\subseteq$ H imply H ∇ G. By Axiom 3A, E $\subseteq$ E$\cup$F and F $\subseteq$ E$\cup$F. So E$\cup$F ∇ G.
8.19 Suppose E$\cup$F can be constructed. Then G ∇ E$\cup$F if and only if there exists an event H such that E $\subseteq$ H, F $\subseteq$ H, and G ∇ H.	If G ∇ E$\cup$F, then E$\cup$F is the requisite H. If E $\subseteq$ H, F $\subseteq$ H, and G ∇ H, then E$\cup$F $\subseteq$ H, and hence G ∇ E$\cup$F by Theorem 8.13.
8.20 Λ ∇ E does not hold for any E.	Λ ∇ E would imply $\Lambda \circ$ E by Axiom 8A, in contradiction of Axiom 5.
8.21 E ∇ Λ if and only if E $\circ$ Λ.	By Axiom 8A, E ∇ Λ implies E $\circ$ Λ. If E $\circ$ Λ, then Axiom 8C, with E for G and Λ for F, together with 8.20, implies that E ∇ Λ.
8.22 If E $\leq$ F and E $\circ$ Λ, then F $\circ$ Λ.	Theorems 8.8 and 8.21

8.23 E\F ∇ G if and only if there is a refinement H of E that is disjoint from F and satisfies H ∇ G.	If E\F ∇ G, then E\F is the requisite H. Going the other way, if H is a refinement of E that is disjoint from F, then H ⊆ E\F. So E\F ∇ G follows from Theorem 8.12.
8.24 G ∇ E\F if and only if [H ⊆ E and H∩F = Λ imply G ∇ H].	We know that H ⊆ E and H∩F = Λ imply H ⊆ E\F. This, together with G ∇ E\F, implies G ∇ H by Theorem 8.13. Going the other way, we know that E\F ⊆ E and (E\F)∩F = Λ. So [H ⊆ E and H∩F = Λ imply G ∇ H] implies G ∇ E\F.

To conclude, we put Axioms 8C and 8D in contrapositive form.

THEOREM	Demonstration
8.25 If E ⊆ G and E ≤ F ≤ G, then E ⊆ F.	Axiom 8C
8.26 If E ⊆ I, F ⊆ I, H ≤ E, and H ≤ F, then H ≤ E∩F.	Axiom 8D

We will use Theorem 8.26 in §12.

9. Happening After

After studying the partial order ⊆ in §1, we introduced the construction E∩F, which is the largest part of E that stands in the relation ⊆ with F. We now proceed analogously for the binary relation ≤, by introducing a construction that represents the largest part of E that stands in the relation ≤ with F.

RULE OF CONSTRUCTION	Explanation
After From events E and F, construct the event E^F.	E^F happens when E happens and F happens at the same time or has already happened.

We call E^F the *happening* of E *after* F. Here "after" is used in a weak sense, to mean "simultaneously or later."

AXIOM	Explanation
9A E^F ⊆ E.	When E^F happens, E happens.
9B E^F ≤ F.	When E^F happens, F happens simultaneously or has already happened.
9C If G ∘ E^F, then G ∘ E or G∩E ∇ F.	In a situation G where E does not happen after F, either E does not happen at all, or else E happens without F having happened yet.
9D If E ⊆ I, F ⊆ I, and $G^{E∪F}$ ∘ Λ, then G^E ∘ Λ or G^F ∘ Λ.	When G happens after E∪F, it happens after E or after F.

Contraposition of Axiom 9C produces the following statement.

THEOREM	Demonstration
9.1 If G ⊆ E and G ≤ F, then G ⊆ E^F.	Axiom 9C

The event E^F is the least upper bound in the partial order ⊆ for the events that refine E and follow F. Theorem 9.1 says that it is an upper bound for these events, and Axioms 9A and 9B imply that there is no smaller upper bound. Speaking less precisely, we say that E is the largest part of E that follows F.

As it turns out, the relation ∇ can be defined in terms of the relation ∘ and the construction After. This is spelled out in Theorem 9.4.

THEOREM	Demonstration
9.2 If E ∇ F, then E ∇ E^F.	Axiom 8B, with E^F for G, and Axiom 9B
9.3 If E ∘ E^F, then E ∇ F.	Axiom 9C, with E for G, and Axiom 1A
9.4 E ∇ F if and only if E ∘ E^F.	Theorems 9.2 and 9.3

Here are a couple of consequences of Theorem 9.4.

THEOREM	Demonstration
9.5 E ≤ F if and only if E = E^F.	Axiom 9A and Theorem 9.4
9.6 If E ⊆ I and F ⊆ I, then G ∘ $G^{E\cap F}$ implies G ∘ G^E or G ∘ G^F.	Axiom 8D and Theorem 9.4

Note the analogy of Theorem 9.6 with Axiom 9D.

We now state conditions for E^F to stand on one side or another of the relations ∇ and ∘ with another event.

THEOREM	Demonstration
9.7 E^F ∇ G if and only if there exists H such that H ⊆ E, H ≤ F, and H ∇ G.	If E^F ∇ G, then E^F is the requisite H, by Axioms 9A and 9B. If H ⊆ E, H ≤ F, then H ⊆ E^F by Theorem 9.1, and this, together with H ∇ G, implies E^F ∇ G by Theorem 8.12.
9.8 G ∇ E^F if and only if H ⊆ E and H ≤ F imply G ∇ H.	If G ∇ E^F, then we can use H ⊆ E^F, which follows from H ⊆ E and H ≤ F by Theorem 9.1, to obtain G ∇ H by Theorem 8.13. If H ⊆ E and H ≤ F imply G ∇ H, then we get G ∇ E^F from Axioms 9A and 9B.
9.9 E^F ∘ G if and only if there exists H such that H ⊆ E, H ≤ F, and H ∘ G.	If E^F ∘ G, then E^F is the requisite H, by Axioms 9A and 9B. If H ⊆ E and H ≤ F, then H ⊆ E^F by Theorem 9.1, and this, together with H ∘ G, implies E^F ∘ G by Theorem 1.13.
9.10 G ∘ E^F if and only if G ∘ E or G∩E ∇ F.	The implication to the right is Axiom 9C. If G ∘ E, then G ∘ E^F by Axiom 9A and Theorem 1.14. If G∩E ∇ F, then G ∘ E^F by Theorem 9.2.

The next theorems explore the strong extensionality of E^F.

THEOREM	Demonstration
9.11 If $E^F \circ G^H$, then $E \circ G$ or $F \circ H$.	By Axiom 9C, $E^F \circ G^H$ implies $E^F \circ G$ or $E^F \cap G \nabla H$. From $E^F \circ G$, we obtain $E \circ G$ by Axiom 9A and Theorem 1.13. From $E^F \cap G \nabla H$, we obtain $E^F \nabla H$ by Theorem 8.12. By Axiom 9B and Theorem 8.9, we then obtain $F \nabla H$. Then, from Axiom 8A, we obtain $F \circ H$.
9.12 If $E \subseteq G$ and $F \subseteq H$, then $E^F \subseteq G^H$.	Contraposition of Theorem 9.11
9.13 If $E \subseteq G$, then $E^F \subseteq G^F$.	Theorem 9.12
9.14 If $F \subseteq H$, then $E^F \subseteq E^H$.	Theorem 9.12
9.15 If $E^F \neq G^H$, then $E \neq G$ or $F \neq H$.	Theorem 9.11
9.16 If $E = G$ and $F = H$, then $E^F = G^H$.	Theorem 9.12
9.17 If $E = G$, then $E^F = G^F$.	Theorem 9.13
9.18 If $F = H$, then $E^F = E^H$.	Theorem 9.14

We now provide a necessary and sufficient condition for $E \cup F \nabla G$ that is more interesting than the one given in Theorem 8.18.

THEOREM	Demonstration
9.19 If $E \cup F \nabla G$, then $E \nabla G$ or $F \nabla G$.	Suppose $E \cup F \nabla G$. By Theorem 9.4, $E \cup F \circ (E \cup F)^G$. By Axiom 3B, $E \circ (E \cup F)^G$ or $F \circ (E \cup F)^G$. If $E \circ (E \cup F)^G$, then $E \circ E^G$ by Axiom 3A and Theorems 9.13 and 1.14, whence $E \nabla G$ by Theorem 9.4. Similarly, if $F \circ (E \cup F)^G$, then $F \nabla G$.
9.20 Suppose $E \cup F$ can be constructed. Then $E \cup F \nabla G$ if and only if [$E \nabla G$ or $F \nabla G$].	If $E \cup F \nabla G$, then $E \nabla G$ or $F \nabla G$ by the preceding theorem. If $E \nabla G$ or $F \nabla G$, then $E \cup F \nabla G$ by Theorem 8.12.

There is an interesting parallel between Theorem 9.20 and Axiom 8D, both of which involve the assumptions $E \subseteq I$ and $F \subseteq I$. Axiom 8D implies that under this assumption, $G \nabla E \cap F$ if and only if $G \nabla E$ or $G \nabla F$. Theorem 9.20 says that under this assumption, $E \cup F \nabla G$ if and only if $E \nabla G$ or $F \nabla G$.

Here are some additional properties of the construction E^F.

THEOREM	Demonstration
9.21 $E \cap F \subseteq E^F$.	Theorem 9.1 with $E \cap F$ for G
9.22 $(E^F)^F = E^F$.	The relation $(E^F)^F \subseteq E^F$ is Axiom 9A with E^F for E. The relation $E^F \subseteq (E^F)^F$ follows from Theorem 9.1 with E^F for G and E.
9.23 If $E \subseteq F$, then $(G^F)^E = G^E$.	The relation $(G^F)^E \subseteq G^E$ follows from Axiom 9A and Theorem 9.13. Using Theorems 9.22, 9.14, and 9.13, we obtain and $G^E = (G^E)^E \subseteq (G^F)^E$.

9.24 If $E \subseteq I$ and $F \subseteq I$, then $E^F = E\cap F$.	By Theorem 9.21, it suffices to show that $E^F \subseteq E\cap F$. We assume $E^F \circ E\cap F$ and derive a contradiction. From $E^F \circ E\cap F$, we obtain $E^F \circ E$ or $E^F \circ F$ by Axiom 2A. But $E^F \circ E$ contradicts Axiom 9A. And when we apply Axiom 8C to $E^F \circ F$, we find that $E^F \triangledown F$ (contradicting Axiom 9B), $F \triangledown I$ (contradicting the assumption $F \subseteq I$), or $E^F \circ I$ (contradicting the assumption $E \subseteq I$).
9.25 $(E\cap F)^G = E^G\cap F^G = E^G\cap F$.	By Theorem 9.13, $(E\cap F)^G \subseteq E^G$ and $(E\cap F)^G \subseteq F^G$. So $(E\cap F)^G \subseteq E^G\cap F^G$ by Theorem 2.1. Because $F^G \subseteq F$ (Axiom 9A), we obtain $E^G\cap F^G \subseteq E^G\cap F$ by Theorem 2.8. So $(E\cap F)^G \subseteq E^G\cap F^G \subseteq E^G\cap F$. From $E^G \subseteq E$ (Axiom 9A again), we obtain $E^G\cap F \subseteq E\cap F$ by Theorem 2.9. By Axiom 9B, $E^G \leq G$, and hence $E^G\cap F \leq G$ by Theorem 8.12. So $E^G\cap F \subseteq (E\cap F)^G$ by Theorem 9.1.
9.26 If $F \subseteq E$, then $F^G = E^G\cap F$.	Theorem 9.25
9.27 If $E \subseteq I$ and $F \subseteq I$, then $(E\cup F)^G = E^G\cup F^G$.	Using the relation $(E\cup F)^G \subseteq E\cup F$, distributivity, and Theorem 9.26, we may write $(E\cup F)^G = (E\cup F)\cap(E\cup F)^G = [E\cap(E\cup F)^G]\cup[F\cap(E\cup F)^G] = E^G\cup F^G$.
9.28 $E^G\backslash F^G = (E\backslash F)^G$.	Using Theorems 6.5 and 9.25, we may write $E^G\backslash F^G = E^G\backslash(E^G\cap F^G) = E^G\backslash(E^G\cap F) = E^G\backslash F$. So we only need to show that $(E\backslash F)^G = E^G\backslash F$. From Theorem 9.13, we obtain $(E\backslash F)^G \subseteq E^G$. We have $(E\backslash F)^G \subseteq E\backslash F$ by Axiom 9A and hence $(E\backslash F)^G\cap F = \Lambda$ by Theorems 6.11 and 5.8. Hence $(E\backslash F)^G \subseteq E^G\backslash F$ by Theorem 6.22. We obtain $E^G\backslash F \subseteq E\backslash F$ from Axiom 9A and Theorem 6.17. And we obtain $E^G\backslash F \leq G$ from Axiom 9B and Theorem 8.12. So $E^G\backslash F \subseteq (E\backslash F)^G$ by Theorem 9.1.
9.29 $E^F = \Lambda$ if and only if [$G \subseteq E$ and $G \leq F$ implies $G = \Lambda$].	If [$G \subseteq E$ and $G \leq F$ implies $G = \Lambda$], then we obtain Axiom $E^F = \Lambda$ by Axioms 9A and 9B. If $E^F = \Lambda$, then we obtain [$G \subseteq E$ and $G \leq F$ implies $G = \Lambda$] by Theorem 9.1.
9.30 If $E \subseteq I$ and $F \subseteq I$, then $H^{E\cup F} = H^E\cup H^F$.	By Theorem 9.14, $H^E\cup H^F \subseteq H^{E\cup F}$. To prove the equality, we set $G = (H^{E\cup F})\backslash(H^E\cup H^F)$. Because $G \leq E\cup F$, $G^{E\cup F} = G$. Using Theorems 9.28, 9.23, and 9.27, we find that $G^E = H^E\backslash(H^E\cup(H^F)^E) = \Lambda$. Similarly, $G^F = \Lambda$. So by Theorem 9.10, $G^{E\cup F} = \Lambda$, or $G = \Lambda$.
9.31 If $E \leq F$ and $E^G = \Lambda$, then $E \leq F\backslash G$.	Since $F = (F\cap G)\cup(F\backslash G)$, Theorem 9.30 yields $E^F = E^{F\cap G} \cup E^{F\backslash G}$. From $E^G = \Lambda$ and Theorem 9.14, we obtain $E^{F\cap G} = \Lambda$, and from $E \leq F$ and Theorem 9.4, we obtain $E = E^F$. So $E^F = E^{F\cap G} \cup E^{F\backslash G}$ reduces to $E = E^{F\backslash G}$. Again using Theorem 9.4, we obtain $E \leq F\backslash G$.

Theorem 9.30 marks our first use of Axiom 9D.

10. Happening Strictly After

Sometimes we are interested in the event that E happens *strictly after* F. This event does not require a new rule of construction; it can be constructed using rules we have already adopted: After and Relative Complement.

CONSTRUCTION	Definition	Reading
$E_2{}^F$	$E^F \backslash F$	The happening of E strictly after F.

We can characterize $E_2{}^F$ order-theoretically: it is the largest refinement of E that requires F but is disjoint from F.

THEOREM	Demonstration
10.1 $E_2{}^F \subseteq E$.	Axiom 9A and Theorem 6.2
10.2 $E_2{}^F \leq F$.	Axiom 9A and Theorems 6.2 and 8.5
10.3 $E_2{}^F \cap F = \Lambda$.	Theorem 6.11
10.4 If $G \subseteq E$, $G \leq F$, and $G \cap F = \Lambda$, then $G \subseteq E_2{}^F$.	By Axiom 9C, $G \subseteq E^F$, and it follows by Theorem 6.22 that $G \subseteq E^F \backslash F$.

Here are some additional properties.

THEOREM	Demonstration
10.5 $E_2{}^F = E^F \backslash (E \cap F)$.	By Theorem 6.5, $E^F \backslash F = E^F \backslash (E^F \cap F)$. By Theorem 9.21, $E^F \cap F = E \cap F$.
10.6 If $E \subseteq I$ and $F \subseteq I$, then $E_2{}^F = \Lambda$.	Theorems 9.22 and 10.5
10.7 If $E \leq F \cup G$ and $E_2{}^F = \Lambda$, then $E \leq F$.	Theorem 9.31

11. Happening Clear

We can also construct, using After and Relative Complement, the event that E happens without F yet having happened. We may call this event the happening of E *clear* of F.

CONSTRUCTION	Definition	Reading
$E \backslash\backslash F$	$E \backslash E^F$	The happening of E without F yet having happened.

Here are some properties of $E \backslash\backslash F$.

THEOREM	Demonstration
11.1 If $G \circ \Lambda$, then $G \nabla F$ or $G \circ E \backslash\backslash F$.	By Theorem 6.34, $G \circ \Lambda$ implies $G \circ G^F$ or lap(G,G^F). By Theorem 9.6, $G \circ G^F$ means $G \nabla F$. On the other hand, lap(G,G^F) means $G^F \circ \Lambda$, which implies, by Axiom 1B, that $G^F \circ G \cap E^F$ or $G \cap E^F \circ \Lambda$. From $G^F \circ G \cap E^F$, we obtain $G^F \circ (G \cap E)^F$ by Theorem 9.25 and then $G \circ E$ by Theorem 9.5. Finally, by Axiom 6A, $G \cap E^F \circ \Lambda$ and $G \circ E$ both imply $G \circ E \backslash E^F$, or $G \circ E \backslash\backslash F$.

11.2 If G ⊆ E\\F and G ≤ F, then G = Λ.	Theorem 11.1
11.3 $(E\backslash\backslash F)^{F}$ = Λ	Theorems 9.22 and 9.28
11.4 $(E\backslash\backslash F)^{F\backslash\backslash E}$ = Λ	Theorems 9.14 and 11.3
11.5 E = (E\\F)∪(E_2^F)∪(E∩F).	By definition, E\\F = E\\E^F, and since E^F ⊆ E, this implies E = (E\\F)∪(E^F). By Theorem 10.5, E_2^F = E^F\\(E∩F), and since E∩F ⊆ E^F, this implies E^F = (E_2^F)∪(E∩F).

12. Diverging and Implying

We say two events *diverge* if they cannot both happen. This is made precise by the following definition.

RELATION	Definition	Reading	Meaning
div(E,F)	If H ≤ E and H ≤ F, then H = Λ.[18]	E and F diverge.	There is no situation (possible event) where both E and F have happened.

This relation is obviously symmetric: div(F,E) if and only if div(E,F). We will use the symmetry without comment. Depending on the context, we will read "div(E,F)" as "E and F are divergent," "E and F diverge," E diverges from F," or "F diverges from E."

The idea of divergence takes on a different significance when we think of one of the instantaneous events involved as a situation. If E is a situation (i.e., if E is an instantaneous event and E ∘ Λ), and div(E,F), then we may say that F has *failed* in the situation E: it is no longer possible for F to happen, because there is no possible situation H after E where F has happened.

If two events cannot both happen, then they cannot both happen at the same time: divergent events are disjoint. This intuition is confirmed by the following theorem.

THEOREM	Demonstration
12.1 If div(E,F), then dis(E,F).	By Axiom 2A and Theorem 8.5, E∩F ≤ E and E∩F ≤ F. So div(E,F) implies E∩F = Λ.

Disjoint events are not necessarily divergent, for one can happen after the other. But disjoint refinements of a single event are divergent.

THEOREM	Demonstration
12.2 If E ⊆ I and F ⊆ I, then div(E,F) if and only if dis(E,F).	Theorems 12.1 and 8.26

[18] Perhaps we should use the constructively stronger condition, "If H ∘ Λ, then H ∇ E or H ∇ F." (In words: in any situation H, either E has not happened or else F has not happened.) It is not clear to us how to choose between the two conditions as the definition of divergence.

Implication is a relation between propositions. As such, it can be regarded as a primitive. It can also be explained semantically: A implies B if B is true in all situations where A is true.

Instantaneous events are not true or false. They happen or fail. So does it make sense to talk about one instantaneous event implying another? What would this mean? A natural way of responding to this question is to shift our attention from the instantaneous events to related propositions. Situating ourselves at the beginning of time, we might consider propositions of the form, "E will happen," where E is an instantaneous event. We could then discuss whether "E will happen" implies "F will happen"—without regard to which happening will come first. Such a discussion seems, however, to require philosophical and mathematical assumptions going far beyond what we have adopted so far.

Instead of speculating about propositions, we will use the concept of divergence, which is already in our framework, to define implication. We can do this because, as we have already noted, the relation div(G,F) can be interpreted by saying that F has failed in the situation G. If we assume that any event eventually happens or fails (it does seem reasonable to make this part of our intuitive concept of an instantaneous event), the proposition that F must happen if E happens is equivalent to the proposition that E must fail if F fails. So we adopt the following definition.

RELATION	Definition	Reading	Meaning
$E \rightarrow F$	If div(G,F), then div(G,E).	E implies F.	Whenever the happening of F is ruled out, the happening of E is also ruled out.

It follows directly from this definition that $\rightarrow$ is reflexive and transitive:

THEOREM
12.3 $E \rightarrow E$
12.4 If $E \rightarrow F$ and $F \rightarrow G$, then $E \rightarrow G$.

If $E \rightarrow F$ and $F \rightarrow E$, we say that E and F are *logically equivalent* instantaneous events. Either both eventually happen or else neither ever happens.

13. Incomparable Events

Let us call two events *incomparable* if neither can happen after the other.

RELATION	Definition	Reading	Meaning
inc(E,F)	$E^F = \Lambda$ and $F^E = \Lambda$.	E and F are incomparable.	Neither E nor F can happen after the other.

The condition $E^F = \Lambda$ means that $G \subseteq E$ and $G \leq F$ implies $G = \Lambda$ (see Theorem 9.29). So the precise meaning of incomparability is that there is no situation in which one of the events happens and the other also happens or has already happened. Divergence is stronger: it says there is no situation in which both events have happened.

In ordinary reasoning, we assume that all events that actually happen fall along a single time-line, and hence incomparability is the same as divergence. If two events both happen, one must happen after or simultaneously with the other. In the theory of relativity, however, there is no universal time-line. In the relativistic world, E can be said to happen after F only if the locations of the two events in space and time permit the news of F's happening, traveling at the speed of light from where F happens to where E happens, to arrive by the time E happens. If the two events are far apart in space, it may be that neither precedes the other in this sense. The two events are then incomparable, and yet if they both happen there is a later situation where they are both in the past (where the news of both has arrived), and hence they are not divergent.

The axioms and rules of construction we have adopted so far appear do not involve assuming the existence of a universal time-line, and they appear, therefore, to be valid for a relativistic world. It is interesting that most of the axioms we need for causal reasoning are valid in this general context. We are primarily interested, however, in systems for ordinary causal reasoning, not in systems for space travelers. So in the next section we will finally adopt axioms that are valid only in a non-relativistic world, where there is a universal time-line.

The following theorems clarify the relationship among the concepts of incomparability, divergence, and disjointness under the axioms adopted so far.

THEOREM	Demonstration
13.1 If div(E,F), then inc(E,F).	Theorem 9.29
13.2 If inc(E,F), then dis(E,F).	Theorem 9.21
13.3 If $E \subseteq I$ and $F \subseteq I$, then div(E,F), inc(E,F), and dis(E,F) are all equivalent.	Theorem 12.2

Here is an additional property of incomparability.

THEOREM	Demonstration
13.4 If $E \subseteq I$ and $F \subseteq I$, inc(E,G), and inc(F,G), then inc($E \cup F$,G).	Theorems 9.27 and 9.30

Here are some examples of incomparable events.

THEOREM	Demonstration
13.5 inc($E_2{}^F$,$F_2{}^E$).	By Theorem 9.29 (and symmetry), it suffices to show that $G \subseteq E_2{}^F$ and $G \le F_2{}^E$ imply $G = \Lambda$. Obtaining $G \subseteq E$ from $G \subseteq E_2{}^F$ and expanding $G \le F_2{}^E$ to $G \le F_2{}^E \le E$, we see, by Theorem 8.25, that $G \subseteq F_2{}^E$, whence $G \subseteq F$. But $G \subseteq E_2{}^F$ also implies that G is disjoint from F. So $G = \Lambda$.
13.6 inc(E\\F,F\\E).	Again, it suffices to show that $G \subseteq$ E\\F and $G \le$ F\\E imply $G = \Lambda$. But $G \le$ F\\E implies $G \le F$, and this implies $G = \Lambda$ by Theorem 11.2.
13.7 inc($E_2{}^F$,$E \cap F$).	This follows from Theorem 13.3, because $E_2{}^F$ and $E \cap F$ are both refinements of E and are disjoint.
13.8 inc(E\\F,$E \cap F$).	This follows from Theorem 13.3, because E\\F and $E \cap F$ are both refinements of E and are disjoint.

Here is an application of Theorem 13.5.

THEOREM	Demonstration
13.9 $E_2^F \leq F \backslash (F_2^E)$.	Theorems 9.31, 10.2, and 13.5

14. Disjoint Merger

We now make our theory non-relativistic by assuming that incomparable events can always be merged into a single instantaneous event.

RULE OF CONSTRUCTION	Explanation
Disjoint Merger From events E and F such that inc(E,F), construct the event E⊕F.	We will adopt axioms that imply E⊕F = E∪F. So the rule says that incomparable events can be merged into a single instantaneous event.

The axioms that we adopt for E⊕F are essentially the same as Axioms 3A and 3B, the axioms we adopted for E∪F in Chapter 3.

AXIOM	Explanation
14A If inc(E,F), authorizing the construction of E⊕F, then E ⊆ E⊕F and F ⊆ E⊕F.	When E happens or F happens, E⊕F happens.
14B If E⊕F ∘ G, then E ∘ G or F ∘ G.	When E⊕F happens, either E happens or F happens.

From Axioms 3A and 3B, we deduced that E∪F, when it can be constructed, is the least upper bound of E and F in the refinement order. We can similarly deduce from Axioms 14A and 14B that E⊕F, when it can be constructed, is also the least upper bound of E and F in the refinement order. It follows that when both E⊕F and E∪F can be constructed, they are equal.

THEOREM	Demonstration
14.1 If inc(E,F), then E⊕F = E∪F.	By the reasoning just explained
14.2 inc(E,F) if and only if div(E,F).	By Theorem 13.1, div(E,F) implies inc(E,F). If inc(E,F), then E ⊆ E⊕F and F ⊆ E⊕F by Axiom 15A, and so div(E,F) by Theorem 13.3.
14.3 E⊕F exists if and only if [E∪F exists and dis(E,F)].	Suppose E∪F exists and dis(E,F). Since E∪F exists, E ⊆ E∪F and F ⊆ E∪F, and hence, by Theorem 13.3, dis(E,F) implies inc(E,F). And so E⊕F exists. Going the other way, if E⊕F exists, then inc(E,F). So E∪F exists by Theorem 14.1, and dis(E,F) by Theorem 13.2.

Theorem 14.3 explains the name *disjoint merger* for E⊕F.

The following theorem clarifies further the relationship between merger and disjoint merger.

THEOREM	Demonstration
14.4 Suppose E, F, and G are pairwise incomparable: inc(E,F), inc(F,G), and inc(E,G). Then (E⊕F)⊕G and E⊕(F⊕G) exist and are equal.	By Theorem 12.1, E⊕F = E∪F. By Theorem 13.4, inc(E∪F,G), and hence (E∪F)⊕G exists. It is equal to both (E⊕F)⊕G and (E∪F)∪G. Analogously, we find that E⊕(F⊕G) exists and is equal to E∪(F∪G).

It follows from this theorem that we can speak unambiguously of the disjoint merger of any number of pairwise incomparable events.

15. The Temporal Lattice

With the help of one additional axiom, we can now show that the temporal order ≤ is a distributive lattice.

We will write E∧F for the greatest lower bound of E and F with respect to ≤, and we will write E∨F for the least upper bound. These constructions are defined as follows.

CONSTRUCTION	Definition	Reading
E∧F	$(E_2{}^F)\oplus(F_2{}^E)\oplus(E\cap F)$	The ending of E and F
E∨F	(E\\F)⊕(F\\E)⊕(E∩F)	The beginning of E and F.

Theorems 13.5 and 13.7 authorize the disjoint merger in the definition of E∧F, while Theorems 13.6 and 13.8 authorize the disjoint merger in the definition of E∨F. We call E∧F the *ending* of E and F because it happens when the happening of the pair ends—either by their happening at the same time or by one happening after the other has already happened. We call E∨F the *beginning* of E and F because the two begin to happen when the first happens. The happening of a pair of events can begin (when one of them happens) without ever ending (because the second never happens).

Here are alternative expressions for E∧F and E∨F.

THEOREM	Demonstration
15.1 $E\wedge F = E^F\cup F^E$.	By Theorems 9.31 and 10.5, $E^F = (F_2{}^E)\cup(E\cap F)$ and $F^E = (E_2{}^F)\cup(E\cap F)$. So $E^F\cup F^E = (E_2{}^F)\cup(F_2{}^E)\cup(E\cap F)$.
15.2 $E\vee F = (E\backslash(E_2{}^F))\cup(F\backslash(F_2{}^E))$.	By Theorem 11.5, $E\backslash(E_2{}^F) = (E\backslash\backslash F)\cup(E\cap F)$ and $F\backslash(F_2{}^E) = (F\backslash\backslash E)\cup(E\cap F)$. So $(E\backslash(E_2{}^F))\cup(F\backslash(F_2{}^E)) = (E\backslash\backslash F)\cup(F\backslash\backslash E)\cup(E\cap F)$.

In order to establish that E∧F and E∨F are the greatest lower bound and least upper bound, respectively, for E and F, we need to prove or assume statements analogous to Axioms 2A, 2B, 3A, and 3B. We adopt as an axiom the statement analogous to Axiom 2B.

AXIOM	Explanation
15 If G ∇ E∧F, then G ∇ E or G ∇ F.	If E and F have not ended happening, at least one of them hasn't happened.

The other statements we prove from axioms already adopted, as follows.

THEOREM	Demonstration
15.3 E∧F ≤ E and E∧F ≤ F.	By Theorem 15.1, we can prove E∧F ≤ E by proving $E^F \cup F^E$ ≤ E. By Axiom 9A, $E^F \subseteq$ E and hence, by Theorem 8.5, E^F ≤ E. By Axiom 9B, F^E ≤ E. So, by Theorem 9.19, $E^F \cup F^E$ ≤ E.
15.4 E ≤ E∨F and F ≤ E∨F.	By Theorem 16.2, we can prove E ≤ E∨F by proving $(E\backslash(E_2{}^F)) \cup E_2{}^F \leq (E\backslash(E_2{}^F)) \cup (F\backslash(F_2{}^E))$. But this follows from applying Theorems 9.19 and 8.15 to Theorem 14.9.
15.5 If E∨F ∇ G, then E ∇ G or F ∇ G.	By Theorem 15.2, E∨F ∇ G can be written $(E\backslash(E_2{}^F)) \cup (F\backslash(F_2{}^E))$ ∇ G. By Theorem 9.19, this implies $E\backslash(E_2{}^F)$ ∇ G or $F\backslash(F_2{}^E)$ ∇ G, whence E ∇ G or F ∇ G.

Using Theorem 15.3 and Axiom 15 just as we used Axioms 2A and 2B in §2, we can establish that E∧F is indeed the greatest lower bound of E and F with respect to ≤, with properties analogous to the properties for E∩F listed in Theorems 2.1 through 2.13. Similarly, using Theorems 15.4 and 15.5, we can establish that E∨F is the least upper bound of E and F with respect to ≤, with properties analogous to the properties for E∪F listed in Theorems 3.1 through 3.13. Because both E∧F and E∨F exist for every pair E and F, we may conclude that the temporal order ≤ is a lattice.

Moreover, as the following theorem establishes, this lattice obeys the resolution axiom.

THEOREM	Demonstration
15.6 If E ≤ F∨G and E∧G ≤ F, then E ≤ F.	By Theorem 11.5, our task is to show that $(E\backslash\backslash G) \cup (E_2{}^G) \cup (E \cap G)$ ≤ F. The condition E∧G ≤ F tells us that $(E_2{}^G) \cup (E \cap G)$ ≤ F. And the condition E ≤ F∨G tells us that E\\G ≤ (F\\G)⊕(G\\F)⊕(F∩G). Since $(F\backslash\backslash G)^G = \Lambda$, it follows by Theorem 10.7 that E\\G ≤ (F\\G)⊕(F∩G), and hence that E\\G ≤ F.

As we learned in §4, this implies that the lattice is distributive.

Here are some further results.

THEOREM	Demonstration
15.7 (E\\F)∧(F\\E) = Λ.	Theorems 11.4 and 15.1
15.8 If E ⊆ I and F ⊆ I, then E∧F = E∩F.	Theorem 10.6

16. Axioms for Event Spaces: Summary

Here is a summary of our relations, constructions, and axioms. Altogether, there are two primitive relations, six primitive constructions, and twenty-one axioms. When a set with the two relations and six constructions obey the twenty-one axioms, we call the set an *event space*.

RELATIONS
$E \circ F$
$E \nabla F$

RULES OF CONSTRUCTION
Overlap From events E and F, construct the event $E \cap F$.
Merger From events E, F, I and proofs of not($E \circ I$) and not($F \circ I$), construct the event $E \cup F$.
Impossible Event Construct the event Λ.
Relative Complement From events E and F, construct the event $E \backslash F$.
After From events E and F, construct the event E^F.
Disjoint Merger From events E and F such that inc(E,F), construct the event $E \oplus F$.

DEFINED PREDICATES	Definition
Possible poss(E)	$E \circ \Lambda$
Impossible imposs(E)	not(poss(E))

DEFINED RELATIONS	Definition
Refines $E \subseteq F$	not($E \circ F$)
Unequal $E \neq F$	$E \circ F$ or $F \circ E$
Equals $E = F$	$E \subseteq F$ and $F \subseteq E$
Requires $E \leq F$	not($E \nabla F$)
Overlaps lap(E,F)	$E \cap F \circ \Lambda$
Disjoint dis(E,F)	$E \cap F = \Lambda$
Diverges div(E,F)	If $H \leq E$ and $H \leq F$, then $H = \Lambda$.
Implies $E \Rightarrow F$	If div(G,F), then div(G,E).
Incomparable inc(E,F)	$E^F = \Lambda$ and $F^E = \Lambda$.

DEFINED CONSTRUCTIONS	Definition
Strictly After E_2^F	$E^F \backslash F$
Clear $E \backslash\backslash F$	$E \backslash E^F$
Ending $E \wedge F$	$(E_2^F) \oplus (F_2^E) \oplus (E \cap F)$
Beginning $E \vee F$	$(E \backslash\backslash F) \oplus (F \backslash\backslash E) \oplus (E \cap F)$

AXIOMS
1A $E \subseteq E$.
1B If $E \circ F$, then $E \circ G$ or $G \circ F$.
2A $E \cap F \subseteq E$ and $E \cap F \subseteq F$.

2B If $G \circ E \cap F$, then $G \circ E$ or $G \circ F$.
3A If $E \subseteq I$ and $F \subseteq I$, then $E \subseteq E \cup F$ and $F \subseteq E \cup F$.
3B If $E \cup F \circ G$, then $E \circ G$ or $F \circ G$.
4 If $E \subseteq F \cup G$ and $E \cap G \subseteq F$, then $E \subseteq F$.
5 $\Lambda \subseteq E$.
6A $E \backslash F \circ G$ if and only if $E \circ (E \cap F) \cup (E \cap G)$.
6B $G \circ E \backslash F$ if and only if $G \circ E$ or $E \cap F \cap G \circ \Lambda$.
8A If $E \nabla F$, then $E \circ F$.
8B If $E \nabla F$, then $E \nabla G$ or $G \nabla F$.
8C If $E \circ F$, then $E \nabla F$ or $F \nabla G$ or $E \circ G$.
8D If $E \subseteq I$ and $F \subseteq I$, then $G \nabla E \cap F$ implies $G \nabla E$ or $G \nabla F$.
9A $E^F \subseteq E$.
9B $E^F \leq F$.
9C If $G \circ E^F$, then $G \circ E$ or $G \cap E \nabla F$.
9D If $E \subseteq I$, $F \subseteq I$, and $G^{E \cup F} \circ \Lambda$, then $G^E \circ \Lambda$ or $G^F \circ \Lambda$.
14A If inc(E,F), authorizing the construction of $E \oplus F$, then $E \subseteq E \oplus F$ and $F \subseteq E \oplus F$.
14B If $E \oplus F \circ G$, then $E \circ G$ or $F \circ G$.
15 If $G \nabla E \wedge F$, then $G \nabla E$ or $G \nabla F$.

The concept of an event space is very general. An event space, like a Boolean algebra, can be finite or infinite. Although our axioms are motivated in part by pictures of finite trees, they do not require that time be discrete, or that Nature foresee only a finite number of possibilities for the situation a short time in the future, or that time end at any point. On the contrary, these axioms are fully consistent with continuous time, in which situations do not necessarily have immediate successors. They permit a situation to decompose into a continuum of alternatives, and they permit open sequences of situations—sequences of situations $E_1, E_2, \ldots$ such that $E_{i+1} \leq E_i$ but there is no E with $E \leq E_i$ for all i. (See Shafer 1998a.)

Like all constructive axiomatizations, our axiomatization is open-ended; we can introduce further constructions if we want. Some possibilities are discussed in the appendix.

IV. Classical Axioms for Event Spaces

We now translate our axioms into classical form and use the classical form to prove a generalization of the Stone representation theorem, as in Shafer (1998a).

Our work in this section brings us back to the intuitive idea with which we began in Section 3 of Part II: an instantaneous event can be represented as a clade in an event tree. We show first that the set of all clades in an event tree satisfies our classical axioms for an event space. Our representation theorem is a converse to this statement: any other structure that satisfies our classical axioms for an event space is isomorphic to a collection of clades in an event tree.

The reasoning in this section is necessarily classical. A clade, by definition, is a subset, which is a classical mathematical idea. And the proof of the representation theorem involves classical reasoning about ultrafilters. But the consequences of our axioms that we cite in the course of our reasoning are all theorems that we proved constructively in Part III.

1. The Classical Axioms

From a classical viewpoint, it is most natural to take $\subseteq$ and $\leq$ as basic and then to have the constructions emerge from the axioms. We begin with a set of objects that we call events, with the assumption that we know what is meant by equality for these objects, and with the assumption that we can substitute equals for equals in any relation. We posit two relations on the set:

RELATION	READING
$E \subseteq F$	E refines F
$E \leq F$	E requires F

We adopt the following axioms for the refinement ordering $\subseteq$. We leave it to the reader to verify that these axioms are classically equivalent to the constructive axioms formulated in §§1-6 of Part III and thus imply that the refinements of a fixed event form a Boolean algebra.

AXIOM	COMMENTS
1 The relation $\subseteq$ is a partial order.	This means that $\subseteq$ is reflexive ($E \subseteq E$), transitive ($E \subseteq F$ and $F \subseteq G$ implies $E \subseteq G$), and antisymmetric ($E \subseteq F$ and $F \subseteq E$ imply $E = F$).
2 Every pair of events E and F have a greatest lower bound.	A greatest lower bound in a partially ordered set is unique. We write $E \cap F$ for the unique greatest lower bound of E and F.
3 If E and F have an upper bound, then they have a least upper bound.	A least upper bound in a partially ordered set is unique. We write $E \cup F$ for the unique least upper bound of E and F when it exists.

4 If $E \subseteq F \cup G$ and $E \cap G \subseteq F$, then $E \subseteq F$.	As observed in §4 of Part III, this implies that the distributive laws hold whenever the least upper bounds in them exist.
5 There exists an event that refines every other event.	This event is unique; we designate it by Λ.
6 For every pair of events E and F, there exists a complement of F relative to E.	A complement of F relative to E is an event H such that $(E \cap F) \cap H = \Lambda$ and $(E \cap F) \cup H = E$. With the help of the preceding axioms, it can be shown that such a complement is unique. We write E\F for the unique complement of F relative to E

Next we adopt axioms for the temporal ordering ≤ that are classically equivalent to the constructive axioms formulated in §§8-9 of Part III.

AXIOM	**COMMENTS**
8A If $E \subseteq F$, then $E \leq F$.	
8B The relation ≤ is a partial order.	
8C If $E \leq F \leq G$ and $E \subseteq G$, then $E \subseteq F$.	
8D If $E \subseteq I$, $F \subseteq I$, $G \leq E$, and $G \leq F$, then $G \leq E \cap F$.	Looking back from the situation G, we may say that if two refinements of an event have happened, they must have happened simultaneously.
9 There is a largest refinement of E that requires F.	This refinement is easily seen to be unique; we designate it by E^F.
9D If $E \subseteq I$, $F \subseteq I$, $G^E = \Lambda$ and $G^F = \Lambda$, then $G^{E \cup F} = \Lambda$.	If G is impossible in both E and F, then G is impossible in $E \cup F$.

Finally, we adopt classical versions of the axioms that are not valid in a relativistic world.

14 If $E^F = \Lambda$ and $F^E = \Lambda$, then E and F have an upper bound.	This means the least upper bound $E \cup F$ exists.
15 $G \leq E$ and $G \leq F$ imply $G \leq E^F \cup F^E$.	The existence of $E^F \cup F^E$ follows from the preceding axiom.

We call a set Ξ an *event space* if has relations $\subseteq$ and ≤ satisfying these axioms. A bijection between two event spaces is an *isomorphism* if it preserves the two relations $\subseteq$ and ≤. A subset Ξ_0 of an event space Ξ is itself an event space with the same relations $\subseteq$ and ≤ provided that (1) $\Lambda \in \Xi_0$, (2) $E \cap F \in \Xi_0$, $E \backslash F \in \Xi_0$, and $E^F \in \Xi_0$ whenever $E \in \Xi_0$ and $F \in \Xi_0$, and (3) $E \cup F \in \Xi_0$ whenever $E \in \Xi_0$, $F \in \Xi_0$, and E and F have an upper bound in Ξ.

2. The Clades in an Event Tree Form an Event Space

As we learned in Part I, an *event tree* is a set $\mathfrak{I}$ with a partial order $\leq^t$ such that two elements S and T are comparable ($S \leq^t T$ or $T \leq^t S$) whenever they have a common lower bound (there exists an element U such that $U \leq^t S$ and $U \leq^t T$). Let us call the elements of an event tree *situations*. A subset E of an event tree is called a *clade* if every two distinct situations in E are incomparable (neither $S \leq^t T$ nor $T \leq^t S$). The empty set qualifies as a clade.

Write $\Xi(\Im)$ for the set of all clades in the event tree $\Im$, and define relations $\subseteq$ and $\leq$ on $\Xi(\Im)$ as follows:

- $E \subseteq F$ means that E is a subset of F.
- $E \leq F$ means that for every $S \in E$ there exists $T \in F$ such that $S \leq^t T$.

We leave it to the reader to verify that this defines an event space—i.e., to check that all the axioms listed in §1 are satisfied. In fact, we carried out this verification, implicitly but thoroughly, when we explained the constructive versions of these axioms in Part III.

3. The $\subseteq$ Ultrafilters in an Event Space Form an Event Tree

Now we show that starting from an event space we can construct an isomorphic space of clades in an event tree. Our demonstration is a straightforward generalization of the usual demonstration of the Stone representation theorem for Boolean algebras (Davey and Priestley, 1990, p. 196). Starting with an arbitrary Boolean algebra, Stone's theorem constructs an isomorphic algebra of sets—the points in the sets being ultrafilters with respect to the partial order in the Boolean algebra. In our generalization, ultrafilters with respect to our partial order $\subseteq$ turn out to be nodes in an event tree, and the event space is then seen to be isomorphic to a set of clades in this event tree. (Ultrafilters with respect to our other partial order, $\leq$, correspond to paths down the tree.)

The concept of an ultrafilter can be defined in any semilattice Ξ with a zero Λ. (A semilattice is a set with a partial order $\subseteq$ in which any two elements E and F have a greatest lower bound $E \cap F$.) We call a nonempty subset S of such a semilattice Ξ an *ultrafilter*[19] if

(1) $\Lambda \notin S$,
(2) if $E \in S$ and $F \in S$, then $E \cap F \in S$, and
(3) if T is a subset of Ξ containing S, $\Lambda \notin T$, and $E \cap F \in T$ whenever $E \in T$ and $F \in T$, then $S = T$.

In other words, an ultrafilter in Ξ is a maximal subset satisfying (1) and (2). Since an ultrafilter is maximal, two distinct ultrafilters must each contain an event not in the other. By the axiom of choice, any nonempty subset S of Ξsatisfying (1) and (2) is contained in an ultrafilter (Davey and Priestley, p. 189). Because the set containing a single element of Ξ satisfies (1) and (2), this implies that each element of Ξ is a contained in an ultrafilter.

Here are some general properties of ultrafilters—properties that hold in any semilattice with a zero.

[19] A more common way of defining this concept begins with the concept of filter. A *filter* is a subset of the semilattice that satisfies (i) if $E \in S$ and $E \subseteq F$, then $F \in S$, and (ii) if $E \in S$ and $F \in S$, then $E \cap F \in S$. A filter S is *proper* if it does not contain the whole semilattice; or, equivalently, if $\Lambda \notin S$. A filter is *maximal* if there is no distinct filter that contains it. We then say that an ultrafilter is a maximal proper filter. Condition (1) is the condition that the ultrafilter be proper. Condition (2) is the same as (ii), and condition (3) is the condition of maximality. Condition (ii) appears to be missing from our definition, but as we shall see (Theorem A2), it is implied by the maximality.

THEOREM	Demonstration
A1. Suppose S is an ultrafilter. Then $E \in S$ if and only if $E \cap F \neq \Lambda$ for all $F \in S$.	Conditions (1) and (2) imply that if $E \in S$, then $E \cap F \neq \Lambda$ for all $F \in S$. On the other hand, if $E \cap F \neq \Lambda$ for all $F \in S$, then $S \cup \{E\} \cup \{E \cap F \mid F \in S\}$ satisfies (1) and (2) and hence is equal, by (3), to S, and hence $E \in S$.
A2. Suppose S is an ultrafilter, $E \in S$, and $E \subseteq F$. Then $F \in S$.	By the preceding theorem, $E \cap G \neq \Lambda$ for all $G \in S$. Because $E \subseteq F$, it follows that $F \cap G \neq \Lambda$ for all $G \in S$. Hence, again by the preceding theorem, $F \in S$.

The following theorems, familiar from the theory of ultrafilters in Boolean algebras, are also valid in the more general context of an event space.

THEOREM	Demonstration
A3. Suppose S is an ultrafilter and $E \in S$. Then $F \in S$ or $E \backslash F \in S$.	Set $T := S \cup \{F\} \cup \{F \cap G \mid G \in S\}$ and $U := S \cup \{E \backslash F\} \cup \{(E \backslash F) \cap G \mid G \in S\}$. Both T and U contain S and are closed under pairwise overlap. At least one of them is composed exclusively of possible events. Indeed, if Λ were in both T and U, then we would have $F \cap G_1 = \Lambda$ and $(E \backslash F) \cap G_2 = \Lambda$, where $G_1 \in S$ and $G_2 \in S$. This would imply $F \cap G = \Lambda$ and $(E \backslash F) \cap G = \Lambda$, where $G = G_1 \cap G_2$, and hence $E \cap G = \Lambda$, contradicting the definition of ultrafilter. So either T or U is an ultrafilter and hence is equal to S.
A4. Suppose S is a nonempty set of events, (1) $\Lambda \notin S$ and (2) if $E \in S$ and $F \in S$, then $E \cap F \in S$, and (3) if $E \in S$, then $F \in S$ or $E \backslash F \in S$. Then S is an ultrafilter.[20]	Let T be an ultrafilter containing S. Suppose $F \in T$. Choose an element E of S. By (3), either $F \in S$ or $E \backslash F \in S$. But $E \backslash F$ cannot be in S, for if it were, $(E \backslash F) \cap F$, which is equal to Λ, would also be in T. So $F \in S$.
A5. $E \subseteq F$ if and only if F is in all ultrafilters that contain E.	Theorem A2 says that if $E \subseteq F$, then F is in all ultrafilters containing E. On the other hand, if $E \subseteq F$ does not hold, then $E \backslash F \neq \Lambda$, and any ultrafilter that contains $E \backslash F$ will also contain E but not F.
A6. A nonempty set of events is an ultrafilter if and only if (1) $\Lambda \notin S$ and (2) if $E \in S$ and $F \in S$, then $E \cap F \in S$, and (3) if $E \in S$, $E_1 \cap E_2 = \Lambda$ and $E_1 \cup E_2 = E$, then $E_1 \in S$ or $E_2 \in S$.	Theorems A3 and A4

[20] This theorem, together with the preceding theorem, shows that the condition that $E \in S$ implies $F \in S$ or $E \backslash F \in S$ can replace the condition of maximality in the definition of an ultrafilter. A proper filter in a Boolean algebra is said to be *prime* when this condition is satisfied, and so in the theory of Boolean algebras Theorem A4 can be expressed by saying that a prime filter is an ultrafilter (Davey and Priestley, p. 187).

Theorem A6 is intuitive if we think of an ultrafilter as representing an indefinitely precise specification of what happens at a particular instant. We call an element E of our event space an *event*; but E can have only some limited degree of detail. If E actually happened, much else happened at the same time. If E_1 and E_2 partition E, then E_1 and E_2 represent more detailed but mutually exclusive accounts of what happened when E happened. The fact that one of these more detailed accounts must always be in S means that S specifies everything that happened at that moment.

The different ultrafilters containing an event E constitute different ways of filling out the details about what else happens along with E. If we write E* for the set consisting of all the ultrafilters that contain E, then E* can be thought of as E in different clothing: E is a partial description of what happened, whereas E* is the set of all complete descriptions that are consistent with this partial description. The equivalence between E and E* is the intuitive content of Stone's representation theorem.

We conclude our survey of the basic properties of ultrafilters in an event space with two theorems that do not appear in the theory of Boolean algebras because they involve our second partial order, ≤.

THEOREM	Demonstration
A7. Suppose $E \le F$. Then for every ultrafilter S containing E, there is an ultrafilter T containing F such that for every $F' \in T$, there is an $E' \in S$ such that $E' \le F'$.	Suppose S is an ultrafilter containing E. Set $T_0 := \{F_0 \mid F_0 \subseteq F \text{ and } E_0 \le F_0 \text{ for some } E_0 \in S\}$. By construction, F is in T_0, and for every $F_0 \in T$, there is an $E_0 \in S$ such that $E_0 \le F_0$. Set $T := \{F_1 \mid F_0 \subseteq F_1 \text{ for some } F_0 \in T_0\}$. It has the same properties. So we only need to show that T is an ultrafilter. We do this using Theorem A6. If $F_1 \in T$, there exists $E_0 \in S$ satisfying $E_0 \le F_1$. Because S is an ultrafilter, $E_0 \ne \Lambda$. So by Theorem 8.22, $F_1 \ne \Lambda$. Now we need to show that $F_1 \cap F_2 \in T$ if $F_1 \in T$ and $F_2 \in T$. It suffices to establish the analogous property for T_0—i.e, to show that $F_1 \cap F_2$ is in T_0 if F_1 and F_2 are in T_0. Choose E_1 and E_2 in S such that $E_1 \le F_1$ and $E_2 \le F_2$. By Theorem 8.14, $E_1 \cap E_2 \le F_1$ and $E_1 \cap E_2 \le F_2$. So by Theorem 8.26, $E_1 \cap E_2 \le F_1 \cap F_2$. Since $E_1 \cap E_2 \in S$, this implies $F_1 \cap F_2 \in T_0$. Now we need to show that $F_1 \in T$ or $F_2 \in T$ whenever $F' \in T$, $F_1 \cap F_2 = \Lambda$, and $F_1 \cup F_2 = F'$. It suffices to establish the analogous property for T_0—i.e, to show that if $F_1 \in T_0$ or $F_2 \in T_0$ whenever $F_0 \in T_0$, $F_1 \cap F_2 = \Lambda$, and $F_1 \cup F_2 = F_0$. Choose $E_0 \in S$ such that $E_0 \le F_0$. Set $E_1 := E_0^{F_1}$ and $E_2 := E_0^{F_2}$. By Theorems 9.5 and 9.30, $E_1 \cup E_2 = E_0$. By Theorem 8.26, $E_1 \cap E_2 \le F_1 \cap F_2$, or $E_1 \cap E_2 \le \Lambda$, or $E_1 \cap E_2 = \Lambda$. Because S is an ultrafilter, $E_1 \in S$ or $E_2 \in S$. So $F_1 \in T_0$ or $F_2 \in T_0$.

A8. Suppose that for every ultrafilter S containing E, there is an ultrafilter T containing F such that for every F' ∈ T, there is an E' ∈ S such that E' ≤ F'. Then E ≤ F.	We argue by contradiction. Suppose E ≤ F does not hold. Then by Theorem 9.2, $E\backslash(E^F) \neq \Lambda$. Choose an ultrafilter S that contains $E\backslash(E^F)$ and hence E. Choose an E' ∈ S such that E' ≤ F. Because S is an ultrafilter, $E' \cap (E\backslash E^F) \in S$. But by Theorem 11.2, $E' \cap (E\backslash E^F) = \Lambda$, contradicting the assumption that S does not contain Λ.

As we know, E ≤ F means that E can happen only after F has already happened. No matter what else happens when E happens (no matter what ultrafilter S containing E we choose), F has already happened in some way (corresponding to some ultrafilter T containing F). Intuitively, S and T are infinitely detailed situations, and T precedes S, because no matter how detailed we make our description F' of T, we can find a description E' of S that is sufficiently detailed to make it clear that F' must have already happened (E' ≤ F'). This is the intuitive content of Theorems A7 and A8.

Now we are ready to show that our ultrafilters form a tree. We write $\Im(\Xi)$ for the set of ultrafilters in Ξ. We define a relation $\leq^t$ on $\Im(\Xi)$ by saying that $S \leq^t T$ if for every F ∈ T there exists E ∈ S satisfying E ≤ F. The next three theorems establish that $\leq^t$ is a partial order.

THEOREM	Demonstration
B1. The relation $\leq^t$ is reflexive: $S \leq^t S$	This follows from the reflexivity of ≤: Because E ≤ E, we may say that for every E ∈ S there exists F ∈ S satisfying E ≤ F.
B2. The relation $\leq^t$ is transitive: If $S \leq^t T$ and $T \leq^t U$, then $S \leq^t U$.	This follows from the transitivity of ≤: For every G ∈ U there exists F ∈ T satisfying F ≤ G and then there exists E ∈ S satisfying E ≤ F and hence, by the transitivity of ≤, E ≤ G.
B3. The relation $\leq^t$ is anti-symmetric: If $S \leq^t T$ and $T \leq^t S$, then S = T.	Suppose $S \leq^t T$ and $T \leq^t S$. We will consider an event E in S and show that it is also in T; by symmetry, this will suffice to prove the theorem. Choose F ∈ T such that F ≤ E. By Theorem A3, either E∩F ∈ T or F\E ∈ T. If E∩F ∈ T, we have our desired conclusion: E ∈ T. So assume F\E ∈ T; we will complete the proof by deriving a contradiction. First choose G ∈ S such that G ≤ F\E. From E∩G ⊆ G and G ≤ F\E we get E∩G ≤ F\E by Theorem 8.14, and from E∩G ≤ F\E ≤ E and E∩G ⊆ E we get E∩G ⊆ F\E by Theorem 8.25, and since E∩G ∈ S, this implies F\E ∈ S. So we obtain E∩(F\E) ∈ S, contradicting the assumption that Λ ∉ S.

Now we show that the partially ordered set $\Im(\Xi)$ is a tree.

B4. Suppose $U \leq^t S$ and $U \leq^t T$. Then E∧F ≠ Λ whenever E ∈ S and F ∈ T.	Consider E ∈ S and F ∈ T. Then there exist $G_1 \in U$ and $G_2 \in U$ such that $G_1 \leq E$ and $G_2 \leq F$. So there exists G ∈ U (namely $G_1 \cap G_2$) such that G ≤ E and G ≤ F. It follows from Axiom 15 that G ≤ E∧F and hence E∧F ≠ Λ.

B5. Suppose S and T are ultrafilters, and E∧F ≠ Λ whenever E ∈ S and F ∈ T. Then for every E ∈ S and F ∈ T, either there exists G ∈ S such that G ≤ F, or there exists G ∈ T such that G ≤ E.	By Theorem A3, either E^F or E\\F is in S, and either F^E or F\\E is in T. We cannot have E\\F ∈ S and F\\E ∈ T, because (E\\F)∧(F\\E) = Λ (Theorem 15.8). So either E^F ∈ S or F^E ∈ T. In the first case, we have an element G of S such that G ≤ F. In the second case, we have an element G of T such that G ≤ E.
B6. Suppose S and T are ultrafilters, and for every E ∈ S and F ∈ T, either there exists G ∈ S such that G ≤ F, or there exists G ∈ T such that G ≤ E. Then $S \leq^t T$ or $T \leq^t S$.	If $S \leq^t T$ fails, then there is some F ∈ T such that there is no G ∈ S satisfying G ≤ F. Then by hypothesis, for every E ∈ S, there exists G ∈ T such that G ≤ E, and hence $T \leq^t S$.
B7. If $U \leq^t S$ and $U \leq^t T$, then $S \leq^t T$ or $T \leq^t S$.	Theorems B4, B5, and B6.

Now we display an isomorphism between our event space Ξ and a set of clades in $\Im(\Xi)$. This is the mapping from E to E*, where

$$E^* := \{S \mid E \in S\}.$$

(Notice that E ∈ S is equivalent to S ∈ E*.) This mapping is one-to-one, because if E ≠ F, then either E\F or F\E is a possible event; if E\F is possible, then there is an ultrafilter containing it, which will be in E* but not in F*, and if F\E is possible, then there is an ultrafilter containing it, which will be in F* but not in E*. We have already shown that the mapping preserves our two partial orders:

- Theorem A5 establishes that E ⊆ F if and only if E* ⊆ F*.
- Theorems A7 and A8 establishes that E ≤ F if and only if E* ≤ F*.

So it is an isomorphism.

V. From Axioms to Logic

In this article, we have provided a concise constructive axiomatization for event spaces, and we have validated this axiomatization in terms of our intuitions about event trees. What use is this? As we explained briefly in the introduction, we have been motivated mainly by the project of using our constructive axioms directly as a logic of events in higher-level type theory. Implementation of this project is beyond the scope of this article, but we need to explain the project in slightly more detail in order to make the significance and value of our work clear. The explanation will be easier if we first discuss the difficulties that arise when we try to use event spaces in a more classical fashion.

1. Putting Events into Classical Logic

Classical (non-intuitionistic) logic makes a clear distinction between syntax and semantics. When we speak of a classical logic, we mean a language with a well-defined syntax—basic symbols together with rules for forming terms and formulae from the basic symbols. Semantics refers to the interpretation of this syntax. In this setup, a mathematical object such as an event space enters primarily as part of the semantics, although its relations and constructions may be mirrored somewhat in the syntax.

Here is one way the classical approach might go in the case of an event space.[21] To construct a language L, we would introduce some symbols for events (including E_L, F_L, and G_L, say), some symbols for relations between events (including $\subseteq_L$ and $\leq_L$), some function symbols ($\cap_L$, $\cup_L$, $\setminus_L$, $\wedge_L$, and $\vee_L$), and so on. We would also introduce symbols for objects of other kinds, for predicates about and relations among these objects, and for the usual logical connectives. Once the syntax is constructed out of all these symbols, we would then speak of models and interpretations. A model would include (1) a set Ξ of events (which should be an event space) and (2) a set Φ of other objects, together with various mathematical predicates, relations, and functions on Φ. An interpretation using such a model would include (3) a mapping of the event symbols and terms constructed from them to elements of Ξ and (4) a mapping of the other object symbols to elements of Φ. If the interpretation maps E_L to the event E in Ξ and F_L to the event F in Ξ, we would require, of course, that it also map $E_L \wedge_L F_L$, a term in L, to the event $E \wedge F$ in Ξ.

This seems unexceptional and not very interesting. What is interesting and leads to difficulties is the fact that the events in Ξ can also be interpreted as situations. Intuitively, whether a certain predicate, say P, holds for a certain object in Φ, say γ, may depend on the situation. This can be incorporated into the formalism in the manner of the situation calculus, which we discuss briefly in Part VI, by adding the situation as another argument of the predicate. What we were calling a predicate symbol in L, say P_L, then becomes a symbol for a relation between objects in Φ and situations in Ξ. Given an event symbol γ_L and an object symbol E_L, the term $P_L(\gamma_L, E_L)$ says, intuitively, that the object named by γ_L has the

[21] Another way of implementing the classical approach, taken by Scherl and Shafer (1998), is to represent actions instead of events in the syntax.

property named by P_L in the situation named by E_L. This term will be mapped to the truth value *True* by an interpretation if $P(\gamma,E)$ holds, where P is the relation to which the interpretation maps the symbol P_L, γ the object to which it maps γ_L, and E the event to which it maps E_L. Otherwise it will be mapped to *False*.

But wait! The last sentence of the preceding paragraph cannot be right, because our situations differ in their specificity. The situation E might specify that the object γ has the property P, and it might specify that γ does not have the property P. But it might also be too broad to say either. It might decompose into two non-zero refinements E_1 and E_2, such that γ has the property P in E_1 but does not have the property P in E_2.

So what shall we do? Shall we say that we have a logic with three truth values—*True*, *False*, and *Maybe*—and forge ahead? This option is explored in Shafer (1998b), but not with clear success. While coherent, a classical three-valued logic of events is not simple and has limited appeal. For many people, the payoff of the superstructure formed by the semantics in classical logic is that it gives a definite meaning to the terms and formulae in the language, and those who take this point of view will have little use for an exceptionally unwieldy semantics that does not even arrive at a definite meaning in many cases.

2. Putting Events into Intuitionistic Logic

The contemporary intuitionistic literature offers a different approach to semantics. In this approach, as articulated by Martin-Löf (1982, 1984), logic is seen as functional programming, and the meaning of terms and propositions is found in the rules for computing with them. The user of the logic establishes the meaning of a symbol by declaring the functional type of the symbol and declaring additional functions for working with it. The concept of assigning a truth-value to a proposition symbol P is replaced by the concept of providing a proof p for the proposition represented by P. The relation between a proposition P and a proof p of it is itself expressed as a type declaration; we say that P is the type of p, and we write

$$p : P.$$

The ordinary logical connectives are brought into type theory by declaring certain functions. For example, conjunction involves a function that forms a new proposition from a pair of propositions, another function that forms proofs of the new proposition from proofs of the pair, and yet another function that yields proofs of the individual propositions from proofs of the pair. In the higher-level version of Martin- Löf's type theory (Nordström et al. 1990, Ranta 1994), all logical judgments are expressed as type declarations.

When a mathematical theory is axiomatized constructively, in the style we followed in Part III, it can be used directly in this logical framework. The rules of construction and axioms of the theory are simply added as additional type declarations. For example, a rule that permits the construction of a line from two points is a function that maps a pair of points to a line. To adopt the rule, we declare the existence of a function of this type; this is the type declaration

$$l : (\text{point})(\text{point})\text{line}.$$

Axioms can also be declared as functions, which supply proofs of certain propositions from proofs of certain other propositions. See the axiomatization of elementary geometry in type theory by von Plato (1995).

The logic of events that we envision begins by adding our axiomatization of event spaces to the higher-level type theory in the style of von Plato. Then we develop a syntax for representing both propositions and events. In classical logic, one applies predicate symbols to object symbols in order to construct terms that can name propositions. In type theory, the role of a predicate is played by a function that maps objects of some type to propositions; for example, the type declaration

isyoung : (Person)Prop

says that "isyoung" maps people to propositions (Ranta 1994). Because we take the meaning of "Bill is young" to be relative to the situation, we add the situation as an argument in this type declaration, obtaining

isyoung : (Situation)(Person)Prop.

In addition to these propositional functions, we also declare event functions. For example, the function

die : (Situation)(Person)Event

maps the situation S and the person X to the event that X dies in or after S. Actions are represented similarly. Other functions, which relate propositions and events in various ways, can be added as the need arises in particular applications.

These ideas can be developed in any computational system that supports higher-level type theory. These include:

- ALF (Magnusson and Nordström 1994), which hews strictly to Martin-Löf's predicative type theory (Nordström, Petersson, and Smith 1990),
- Nuprl (Constable 1986), based on an older variant of Martin-Löf's theory,
- Coq (Dowek et al. 1993), based on Coquand and Huet's calculus of constructions (Coquand and Huet 1988), and
- Lego (Luo and Pollack 1992), based on an extended version of the calculus of constructions (Luo 1994).

Isabelle (Paulson 1994) is a more generic logical framework; it supports a variety of logics, including classical first-order logic and Zermelo-Fraenkel set theory as well as constructive logics based on type theory. These systems are all completely and provably adequate as proof-checkers. They vary in the extent to which they provide tactics and other facilities for theorem proving, but of course they provide only starting points for the development of practical reasoning systems in any particular domain. The development of a practical version of our logic of events within one these systems must therefore be regarded as a long-term research project.

3. Conclusion and Future Prospects

In this article, we have provided simpler and more transparent axioms for event spaces. These axioms are constructive in the intuitionistic sense, which means that they can be used as the starting point for computational implementations of event spaces for causal reasoning in specific domains. We believe that this approach to causal reasoning offers greater prospects for the implementation of event-space logics than the approaches previously developed by Scherl and Shafer (1998) and Shafer (1998b), which stay closer to first-order logic. In order to demonstrate the value of our approach, we plan to implement these axioms in one or more of the logical frameworks described in §2. We also anticipate that practical applications of our approach will require extensions to include probabilities.

VI. Comparisons

We conclude with some comments on other well-known approaches to temporal and causal reasoning, which may help readers form their own understanding of how our ideas fit into a wider context.

1. Temporal Logic

Temporal logics augment the syntax of classical logic with operators that represent temporal notions. From a proposition P, these operators allow us to form additional propositions, such as sometimes(P), always(P), nexttime(P), etc. The semantics of a temporal logic always includes an ordered set of time points, or situations, that are used to interpret the temporal operators. In the case of linear temporal logics, these situations fall along a single line, while in branching-time logics, they form a tree. These situations appear only in the semantics, however; the syntax of the language does not provide any way of talking directly about situations, events, or points in time.

Event spaces go beyond the semantic picture used by branching-time temporal logics in at least two ways. First, event spaces allow continuous time, whereas those temporal logics that use the operator "nexttime" require discreteness in the set of situations. Second, and more fundamentally, the concept of refinement in event spaces means that a situation is more than a point in time. Even after a situation is refined enough to specify the time exactly, it can gain further substantive content by further refinement.

The logic of events that we outlined in the preceding section involves facilities for talking explicitly about events and situations, and this is not usually desired in a temporal logic. It is interesting to note, however, that objects functionally equivalent to events tends to emerge in the syntax when branching-time temporal logic is fully developed. This is illustrated by CTL* (Emerson and Halpern 1986; Emerson 1990), a well-known and highly expressive branching-time temporal logic. As it turns out, CLT* involves two different concepts of proposition. On the one hand, the logic has *state formulae*, which are true or false in situations and thus represent propositions of the usual kind. On the other hand, it has *path formulae*, which are true or false only with respect to paths. Intuitively, a path formula is an assertion in future tense. It says that some event will eventually happen, without saying when. Having such a statement in one's language is functionally equivalent to having a name for the event in one's language.

For additional information on temporal logic, see Bolc and Szalas (1995), Goldblatt (1992), or Vila (1994).

2. The Situation Calculus

There is a great deal of work on reasoning about action in artificial intelligence that comes closer to our ideas, because it puts situations (or at least time) into the syntax. This includes work on causal reasoning with explicit time (Shoham 1988, Stein and Morgenstern

1994), as well as substantial body of work on the situation calculus (McCarthy and Hayes 1969, Levesque et al. 1997, Pinto 1994).

As formulated by Reiter (1991), the situation calculus is a sorted first-order language. It has a constant S_0, which denotes the initial situation, and a distinguished binary function symbol "do," which allows us to talk about how actions change the situation. When α denotes an action and s denotes a situation, do(α,s) denotes the situation that results when one performs α in s. Time is not represented explicitly, but time passes as actions are preformed. (For an extension in which time is represented explicitly, see Pinto 1998.) A number of authors (Reiter 1993, Lin and Reiter 1994, Shanahan 1998, Pirri and Reiter 1999, Reiter 2000) have explored foundational axioms for the situation calculus. Shanahan calls them axioms of arboreality, because they require that the space of situations form a tree. One of these axioms is a second-order sentence that rules out the existence of situations that cannot be obtained by actions we have named starting with S_0. The axioms ensure that two situations will be the same if they result from the same sequence of actions applied to the initial situation.

The most fundamental difference between the situation calculus and our event-space approach is that in our approach we can name situations at varying levels of detail. In the situation calculus, the initial situation S_0 is supposed to specify the state of the world in complete detail, at least in all the detail that will ever be needed in the discussion that ensues, and this completeness is supposed to persist as actions are applied. Another artificial intelligence language, the event calculus (Shanahan 1997, Kowalski and Sergot 1986), has been developed to represent partial information about events. But unlike the situation calculus, the event calculus is a linear logic rather than a branching-time logic; it cannot handle alternative sequences of events. Our event-space approach can be seen as a way of accomplishing the purposes of both the situation calculus and the event calculus in a single framework.

Both the situation calculus and the event calculus make assumptions of inertia (Sandewall 1994), which ensure that properties are persistent in the absence of the specified actions. Such persistence is needed in order to plan and reason about actions. Although it does not seem to make sense for us to adopt the situation-calculus axioms of inertia wholesale, we obviously need to deal with the issue of persistence in our logic of events.

Appendix. Additional Constructions

The constructive axiomatization we developed in Part III is open; additional rules of construction, together with additional axioms, can always be added. In particular, we can regain the fuller system of Shafer (1998a) with a few additions. These are reviewed here.

1. Decomposing E with Respect to F

It is most natural, perhaps, to introduce as additional constructions the part of E that diverges from F and the part of E that implies F. We write E_5^F for the part of E that diverges from F—i.e., the largest refinement of E that diverges from F. And we write $E_{\rightarrow}^F$ for the part of E that implies F—i.e., the largest refinement of E that implies F. These constructions are illustrated in Figure 17.

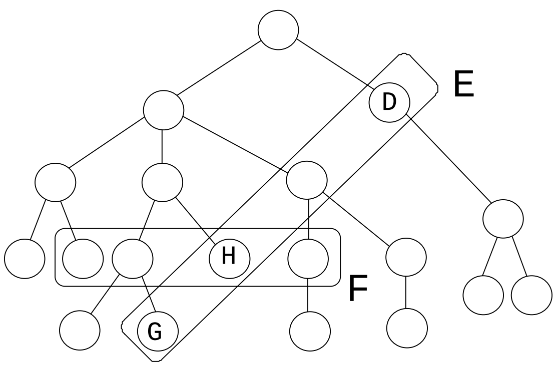

$E_5^F = \{D\}$

$E_{\rightarrow}^F = \{G,H\}$

Figure 17 The parts of E diverging from and implying F.

Once we have defined E_5^F and $E_{\rightarrow}^F$, we can decompose E into five events E_1^F, E_2^F, E_3^F, E_4^F, and E_5^F as follows.

- $E_1^F := E \cap F$. Whenever E_1^F happens, F happens simultaneously.
- $E_2^F := (E^F) \backslash (E \cap F)$. Whenever E_2^F happens, F has already happened strictly earlier.
- $E_3^F := (E_{\rightarrow}^F) \backslash (E^F)$. Whenever E_3^F happens, F is inevitable; it must happen later.
- $E_4^F := E \backslash ((E_1^F) \cup (F_2^E) \cup (E_3^F) \cup (E_5^F))$. Whenever E_4^F happens, F is possible but not inevitable; it may happen later and it may fail later.
- E_5^F we have constructed directly. Whenever E_5^F happens, F is impossible; either it was already impossible or becomes impossible (fails) with the happening of E_5^F.

Some of these five events may sometimes be impossible. But they always decompose E, inasmuch as they all refine E, they do not overlap; and their merger is all of E. Figure 18 gives an example in which each E_i^F is represented by a single node.

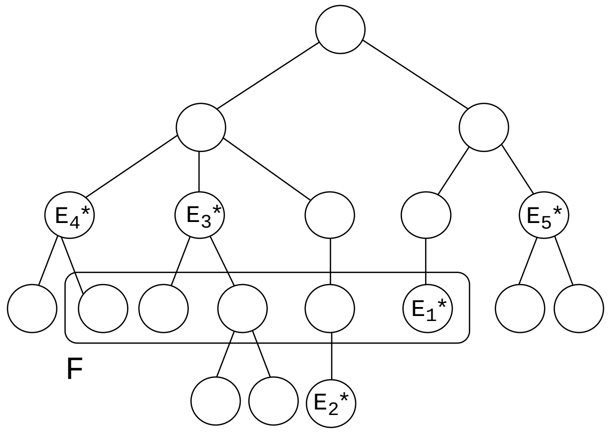

Figure 18 The decomposition of E with respect to F, where E = {E_1*,E_2*,E_3*,E_4*,E_5*}.

Shafer (1998a) bases his axiomatization of event spaces on the constructions E_1^F, E_2^F, E_3^F, E_4^F, and E_5^F, together with relations, E 1 F, E 2 F, E 3 F, E 4 F, and E 5 F, where E i F holds if E_i^F is possible. Using these five constructions and five relations as primitives, he constructs a wide variety of causal relations, which we will not review here.

2. Failure

Figure 19 illustrates the last construction we will consider: the *failure* of E, which we designate E^-. This is the event that happens when E fails—i.e., when Nature (who sees everything that happens as it happens) sees that E has become impossible.

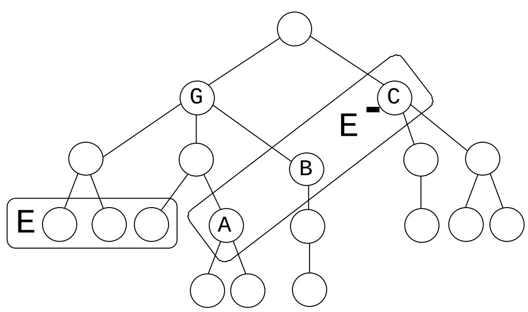

Figure 19 The failure of E. Here we assume that A, B, and C mark the first point in time where it is determined that E will not happen. (This assumption is not part of the meaning of an event tree that we generally give to an event tree.)

Failure is a very powerful construction. We can use it to define the constructions introduced in the preceding section: $E_5^F = E^{(F^-)}$, and $E_{\rightarrow}^F = E^{(F^+)}$, where $E^+ = (E^-)^-$. It is not clear to us, however, that failure should play a fundamental role in casual reasoning. In ordinary causal talk, we often say that an event E makes an event F impossible: Bill's going to the bar made it impossible for him to get home for dinner on time, the 10 inches of snow made it impossible for the plane to take off, etc. But usually we mean only that by the time E happens, F is impossible; we are not very concerned about exactly when F becomes impossible. The plane's take-off may have been blocked before the snowfall total reached ten inches. Pinpointing just when this was may be difficult and not particularly helpful. We have therefore avoided making the construction E^- fundamental in our axiom system.

The construction also has other aspects that are restrictive. For example, it forces us to assume that our event tree has an initial node—a starting point at the top. This assumption is not actually made in the rest of our axiomatic system; all our other axioms and constructions permit the past to be infinite, without beginning. In order to avoid assuming there is an initial situation, Shafer (1998a) used a relative rather than an absolute concept of failure. Instead of assuming that we can construct the failure of E in an absolute sense, he assumed only that we can construct the failure of E after another event F.

Another aspect of failure that makes its representation in event trees confusing, even if we do not assume an initial situation, is that an event tree detailed enough to represent E is not necessarily detailed enough to represent E^{-}. Figure 20 illustrates the point. This figure differs from Figure 19 merely by interpolating an additional situation N between two situations. In general, we do not consider such interpolation a falsification of an event tree, because a tree is never more than a partial description of the possibilities in nature. But the interpolation of N does falsify the additional assumption, made when we identify {A,B,C} as the failure of E, that there are no situations preceding any of them in which Nature already knows E to have failed.

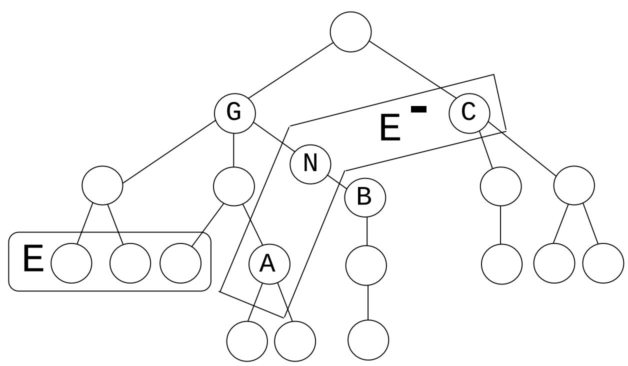

Figure 20 This tree differs from the tree in Figure 19 only in that a node N has been inserted between G and B. The presence of N rules out B's marking the first point in time where Nature knows E will not happen, because Nature already knows this in N. If this new tree is detailed enough to depict E^{-}, then E^{-} is equal to {A,N,C}, not to {A,B,C}. If there are yet other situations, preceding A, N, or C but not shown in this figure, where Nature already knows that E will not happen, then even this figure is not detailed enough to allow the depiction of E^{-}.

From an abstract point of view, the existence of the failure E^{-} is a continuity condition on the partial order $\leq$. It says that the situations in which E has failed have a least upper bound in our space of events, and that E has also failed in this least upper bound. We consider it an open question whether this kind of continuity condition is appropriate or needed for causal reasoning. We are much more comfortable with analogous continuity conditions in the partial order $\subseteq$, because the requisite least upper bounds necessarily exist if we have a representation in terms of an event tree. In an event tree, refinements are represented by subsets, and the algebra of subsets of a set is a complete Boolean algebra: every collection of subsets has a least upper bound.

The fact that the axiomatization in this article does not rely on a concept of failure is one of the major advances of this axiomatization over the one given by Shafer (1998a).

References

Baral, C., M. Gelfond and A. Provetti (1997). Representing actions: Laws, observations and hypotheses. *Journal of Logic Programming* **31** 201-243.

Baras, Bruno et al. (1999) *The Coq Proof Assistant Manual. Version 6.2.4.* January 28, 1999.

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

Bolc, Leonard, and Andrzej Szalas (1995). *Time and Logic: A Computational Approach.* London: UCL Press.

Bru, Bernard (1986). Postface. In Laplace (1986).

Coquand, Thierry, and Gérard Huet (1988). The calculus of constructions. *Information and Computation* **76** 95-120.

Constable, R.L., et al. (1986). *Implementing Mathematics with the NuPRL Proof Development System*. Englewood Cliffs, New Jersey: Prenctice-Hall.

van Dalen, Dirk (1997). *Logic and Structure*. New York: Springer-Verlag.

van Dalen, Dirk (1986). Intuitionistic logic. Pp. 225-339 of *Handbook of Philosophical Logic*, Volume III, edited by D. Gabbay and F. Guenther. Reidel.

Davey, B.A., and H.A. Priestley (1990*). Introduction to Lattices and Order*. New York: Cambridge.

Dowek, G, A. Felty, H. Herbelin, G. Huet, C. Murthy, C. Parent, C. Paulin-Mohring, and B. Werner (1993). *The Coq Proof Assistant User's Guide, version 5.8.* Lyon: INRIA-Rocquencourt and CNRS-ENS. http://pauillac.inria.fr/coq/assis-eng.html

Dummett, Michael (1977). *Elements of Intuitionism*. Oxford: Oxford University Press.

Emerson, E. Allen, and Joseph Y. Halpern (1986). "Sometimes" and "not never" revisited: On branching versus linear time temporal logic. *Journal of the Association for Computing Machinery* **33:1** 151-1711.

Emerson, E. Allen (1990). Temporal and modal logic. Pp. 995-1072 of *Handbook of Theoretical Computer Science,* Vol. B, edited by J. Leeuwen. Amsterdam: Elsevier.

Gelfond, M., and V. Lifschitz (1993). Representing actions and change by logic programs. *Journal of Logic Programming* **17** 301-323.

Goldblatt, Robert (1992). *Logics of Time and Computation*, 2nd Edition, Revised and Expanded. Stanford: CSLI Publications

Kowalski, R.A., and M.J. Sergot (1986). A logic-based calculus of events. *New Generation Computing* **4** 67-95.

Laplace, Pierre Simon de (1986). *Essai philosophique sur les probabilités*. Paris: Christian Bourgois. (Laplace's *Essai* was first published in 1814, and the fifth and definitive edition was published in 1825.)

Levesque, Hector, Raymond Reiter, Yves Lesperance, Fangzhen Lin, and Richard B. Scherl (1997). GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming*.

Lin, Fangzhen, and Raymond Reiter (1997). How to progress a database. *Artificial Intelligence* **92** 131-167.

Lin, Fangzhen, and Raymond Reiter (1994). State constraints revisited. *Journal of Logic and Computation* **4** 566-678.

Luo, Zhaohui (1994). *Computation and Reasoning: A Type Theory for Computer Science.* Oxford: Clarendon Press.

Luo, Zhaohui, and Randy Pollack (1992). Lego Proof Development System: User's Manual. LFCS Report ECS-LFCS-92-211, Department of Computer Science, University of Edinburgh. May 7, 1992. http://www.dcs.ed.ac.uk/home/lego

Magnusson, Lena, and Bengt Nordström (1994). The ALF proof editor and its proof engine. Pp. 213-237 of *Types for Proofs and Programs*. Henk Barendregt and Tobias Nipkow, eds. Lecture Notes in Computer Science No. 806. New York: Springer-Verlag.

Martin, Thierry (1996). *Probabilités et critique philosophique selon Cournot.* Paris: Vrin.

Martin-Löf, Per (1982). Constructive mathematics and computer programming. Pp. 153-175 of *Proceedings of the Sixth International Congress of Logic, Methodology, and Philosophy of Science,* L.J. Cohen, J. Los, H. Pfeiffer, and K.-P. Podewski, eds. Amsterdam: North Holland.

Martin-Löf, Per (1984). *Intuitionistic Type Theory*. Naples: Bibliopolis.

McCarthy, John, and Patrick Hayes (1969). Some philosophical problems from the standpoint of artificial intelligence. Pp. 463-502 of *Machine Intelligence* **4**, edited by B. Meltzer and D. Michie. Edinburgh: University Press.

Nordström, Brengt, Kent Petersson, and Jan M. Smith (1990) *Programming in Martin-Löf's Type Theory*. Oxford: Clarendon Press.

Paulson, L. (1994). *Isabelle: A generic theorem prover.* Lecture Notes in Computer Science 828. New York: Springer. http://www.cl.cam.ac.uk/Research/HVG/Isabelle/dist/

Pinto, Javier A. (1994). Temporal reasoning in the situation calculus. Ph.D. Dissertation, Department of Computer Science, University of Toronto, Toronto, Ontario. Available as report KRR-TR-94-1.

Pinto, Javier A. (1998). Occurrences and narratives as constraints in the branching structure of the Situation calculus. *Journal of Logic and Computation* **9** 777-808.

Pirri, Fioara and Ray Reiter (1999). Some contributions to the metatheory of the situation calculus. *Journal of the ACM* **45** 261-325.

von Plato, Jan (1995). The axioms of constructive geometry. *Annals of Pure and Applied Logic* **76** 169-200.

von Plato, Jan (1996). Organization and development of a constructive axiomatization. Pp. 288-296 of *Types for Proofs and Programs*, edited by S. Berardi and M. Coppo, Lecture Notes in Computer Science, 1158. Springer-Verlag.

von Plato, Jan (1997). Positive partial order, lattices, and Heyting algebras. Unpublished.

Ranta, Aarne (1994). *Type Theoretical Grammar*. Oxford University Press.

Reiter, Raymond (1991). The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. Pp. 359-380 *of Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*. San Diego: Academic Press.

Reiter, Raymond (1993). Proving properties of states in the situation calculus. *Artificial Intelligence* **64** 337-351.

Reiter, Raymond (2000). *Knowledge in Action: Logical Foundations for Describing and Implementing Dynamical Systems*. Book Draft.

Robinson, J.A. (1965). A machine-oriented logic based on the resolution principle. *Journal of the ACM* **12** 23-41.

Sandewall, Erik (1994). *Features and Fluents: The Representation of Knowledge about Dynamical Systems. Volume 1*. Clarendon Press, Oxford.

Scherl, Richard, and Glenn Shafer (1998). A logic of action, causality, and the temporal relations of events. Pp. 89-96 of the *Proceedings of the Fifth International Workshop on Temporal Representation and Reasoning (Time-98),* edited by Lina Khatib and Robert Morris. Los Alamitos, Calif.: IEEE Computer Society.

Shafer, Glenn (1996). *The Art of Causal Conjecture*. Cambridge, Massachusetts: MIT Press.

Shafer, Glenn (1998a). Mathematical foundations for probability and causality. Pp. 207-270 of *Mathematical Aspects of Artificial Intelligence*, edited by Frederick Hoffman. American Mathematical Society, Symposia in Applied Mathematics, Volume 55.

Shafer, Glenn (1998b). Causal logic. *Proceedings of ECAI98 (13th European Conference on Artificial Intelligence),* edited by Henri Prade. Pp. 711-719. Chichester: Wiley.

Shanahan, Murray (1997). *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia.* Cambridge, Massachusetts: MIT Press.

Shoham, Yoav (1988). *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence.* Cambridge, Massachusetts: MIT Press.

Stein, L., and L. Morgenstern (1994). Motivated action theory: a formal theory of causal reasoning. *Artificial Intelligence* **71** 1-42.

Szász, Gábor (1963). *Introduction to Lattice Theory.* Third revised and enlarged edition. New York: Academic Press.

Troelstra, A., and van Dalen, D. (1988). *Constructivism in Mathematics. Volume 2.* Amsterdam: North-Holland.

Vila, L. (1994). A survey on temporal reasoning in artificial intelligence. *AICOM (Artificial Intelligence Communications)* **7:1** 4-28.