

Probability Theory and Causation: A Branching Space-Times Analysis

Thomas Müller

ABSTRACT

We provide a formally rigorous framework for integrating singular causation, as understood by Nuel Belnap's theory of *causae causantes*, and objective single case probabilities. The central notion is that of a causal probability space whose sample space consists of causal alternatives. Such a probability space is generally not isomorphic to a product space. We give a causally motivated statement of the Markov condition and an analysis of the concept of screening off.

- 1 *Causal dependencies and probabilities*
 - 1.1 *Background: causation in branching space-times*
 - 1.2 *What are probabilities defined for?*
- 2 *Basic transitions*
 - 2.1 *Basics of basic transitions*
 - 2.2 *Sets of basic transitions*
- 3 *Causal probability theory*
 - 3.1 *Some simple cases*
 - 3.2 *General causal probabilities*
 - 3.3 *Application: probability of suprema of a chain*

1 Causal dependencies and probabilities

Probability theory describes the likelihood of outcomes of chance set-ups. In this paper, we will assume that the framework of probability theory applies to single, concrete chance set-ups, such as a single coin toss at a specific place and time. This assumption appears to be uncontroversial in view of the fact that probability theory is actually successfully applied. Since we are after a metaphysical picture combining causation and probabilities, we will assume that both causal relations and single case probabilities are objective. The aim

of this paper may be described as trying to make sense of that assumption by showing how it might be true.

Probability theory by itself has nothing to say about causation. However, questions of probabilities and questions of causation are obviously related. For example, one may ask whether probabilities can be a guide to discovering causal relations. The literature on this question is large. When moved from an epistemic to a metaphysical level, the question gets a different reading: What are the objective influences of objectively existing causal relations on objectively existing probabilities? If a concrete situation involves more than one chance set-up, the causal interrelation of the set-ups may impose constraints on the probabilistic structure of the situation. For example, if depending on the outcome of one chance set-up, another one is enabled or not, that causal structure should be mirrored in the probabilistic structure.

In this paper, we will provide a rigorous framework for bringing together causation and probabilities. The paper is based on an existing theory of causation Nuel Belnap's theory of *causae causantes* ([Belnap [2005]], the background for which is *branching space-times* ([Belnap [1992]]). In this paper, we presuppose the notation of Belnap ([2005]), which should be consulted. A very readable introduction to the problem of bringing together causation and probabilities is given in Weiner and Belnap ([2005]), which should also be consulted. The present paper may be read as a sequel to that paper, developing its ideas in more formal detail.

1.1 Background: causation in branching space-times

Branching space-times is a rigorous framework for combining relativistic space-times with indeterminism. The basic building blocks of branching space-times are a set W of concrete point events and a partial order \leq defined on W ; the companion strict order is written ' $<$ '. Histories h are maximal directed sets, i.e., maximal such that for any $e_1, e_2 \in h$ there is $e_3 \in h$ with $e_1 < e_3$ and $e_2 < e_3$. Points e_1, e_2 that are incomparable (neither $e_1 \leq e_2$ nor $e_2 < e_1$) can still be elements of the same history, viz., if they have a common upper bound. In that case, we call e_1 and e_2 *space-like related*, and we write e_1 SLR e_2 . Branching occurs at single indeterministic points. Thus, any two histories h_1, h_2 to which an indeterministic point e belongs are either divided at e ($h_1 \perp_e h_2$) or undivided at e ($h_1 \equiv_e h_2$), the latter being an equivalence relation. The set of histories containing e is written $H_{(e)}$. $H_{(e)}$ is partitioned by \equiv_e ; we write Π_e for the partition and $\Pi_e \langle h \rangle$ for the unique element of Π_e containing h (presupposing $e \in h$). Thus, $\Pi_e \langle h \rangle$ is the set of histories containing e that are undivided from h at e . An *initial* I is a set of point events

that is part of one history. We write $H_{[I]}$ for the set of histories of which I is a part; $H_{[I]} \neq \emptyset$. It is often adequate to think of I as an upper-bounded *initial chain*, and $I = \{e\}$ is an important special case. (A chain is a set of pairwise comparable point events.) The *primary outcomes* of an initial I are the elements of the partition Π_I induced by undividedness at I , where $h \equiv_I h'$ iff $h \equiv_e h'$ for all $e \in I$. An *outcome event* O is a set of (lower-bounded) outcome chains, all of which begin in one history. We write $H_{\langle O \rangle}$ for the set of histories in which O begins to occur (even though it does not have to occur as a whole); by definition, $H_{\langle O \rangle} \neq \emptyset$.

Subscribing to the rather uncontroversial assumption that causation is a two-place relation between cause and effect, the following two questions are crucial for any theory of causation:

1. What is caused? Which entities are *effects*?
2. What causes? Which entities are *causes*?

As is well known, there are numerous different answers to these questions. One main divide is between singular and generic causes and effects. Obviously, both notions are important – consider ‘he broke the vase’ vs. ‘smoking causes cancer’. It seems both ontologically most promising and most in line with our framework to address questions of singular causation first. Thus we look for singular, concrete causes and effects. What sorts of entities are these? A good intuition seems to be that causation involves change, but it is notoriously difficult to say what a change is. The notion of a *transition*, introduced by von Wright ([1963]), is a liberalised notion of change: a transition is ‘something and then something’ not necessarily ‘... something else’. Following von Wright, in causation in branching space-times the crucial causal notion is therefore that of a *transition*. Technically, a transition, written $I \mapsto O$, is just an ordered pair of sets I and O of point events, where I is an initial event and O is an outcome event located after I (i.e., for all $e_I \in I$ and for all $e_O \in O$ we have $e_I < e_O$).

The answer to question (1) above, from the point of view of causation in branching space-times, is accordingly:

1. *Transitions* are caused.

In sharp contrast to other theories of causation, causation in branching space-times assumes that non-trivial causation depends on indeterminism. A transition $I \mapsto O$ whose outcome O is bound to occur, given that I occurs, has *no causes*. It does not need any, since it happens anyway. Causes are thus understood as *originating causes* (*causae causantes*), which seems to capture our actual usage of the notion of causation much more nicely than other accounts.

What are the causes? Causation in branching space-times gives the following answer to question (2) above:

2. Causes are sets of (especially simple) transitions.

Thus, a transition is caused by (possibly a huge number of) transitions. The transitions that figure as causes are of an especially simple sort: they are from a single point e as initial (a *past cause-like locus* of the caused transition) to one of e 's primary outcomes, $\Pi_e\langle h \rangle$. Given an outcome event O of e , $e < O$, the set of histories $\Pi_e\langle O \rangle$ is the uniquely determined member of the partition Π_e of $H_{(e)}$ in which O is kept possible, i.e., for which $\Pi_e\langle O \rangle \cap H_{(O)} \neq \emptyset$. In fact, we then have $H_{(O)} \subseteq \Pi_e\langle O \rangle$. The corresponding transitions, written $t_i = e_i \mapsto H_i$, $H_i \in \Pi_e$, are called *basic transitions*. It turns out, as it should, that basic transitions are their own causes. Thus, they may figure as irreducibly indeterministic building blocks of our indeterministic world. In the theory that will be developed in this paper, they will also form the building blocks for the construction of probability spaces, thus providing a close link between causation and probabilities.

Having mentioned the concepts, we now give the formal definitions of 'past cause-like locus' and 'causa causans' (cf. Belnap [2005]) for later reference.

Definition 1 (Past cause-like loci) Let $I \mapsto O$ be a transition. Then the set of past cause-like loci for that transition, $pcl(I \mapsto O)$, is defined to be

$$pcl(I \mapsto O) := \{e \in W \mid e < O \text{ and } \exists h \in H_{[I]} h \perp_e H_{(O)}\}$$

Thus, a point event e is a past cause-like locus for $I \mapsto O$ iff¹ (i) $e < O$, i.e., e is a lower bound for the outcome O , and (ii) there is a history h in which the initial I finishes ($h \in H_{[I]}$) that is separated at e from all histories in which O begins to occur ($h \perp_e H_{(O)}$). Past cause-like loci make a difference as to the occurrence of O , even given the occurrence of I . They are therefore necessarily indeterministic points, i.e., points with more than one immediate outcome.

Definition 2 (Causae causantes) Let $I \mapsto O$ be a transition. Then the set of causae causantes for that transition, $cc(I \mapsto O)$, is defined to be

$$cc(I \mapsto O) := \{e \mapsto \Pi_e\langle O \rangle \mid e \in pcl(I \mapsto O)\}$$

A *causa causans* for a transition $I \mapsto O$ is thus a transition from a past cause-like locus e to that unique immediate outcome $\Pi_e\langle O \rangle$ of e that keeps O possible right after e .

¹ 'Iff' is short for 'if and only if'.

1.2 What are probabilities defined for?

Intuitively, probabilities are probabilities of outcomes of chance set-ups. Causation in branching space-times offers the notion of an indeterministic transition to stand for ‘chance set-up’. If $I \rightsquigarrow O$ can occur, but O does not have to occur given I , then it will in many cases be sensible to ask *how probable* the transition $I \rightsquigarrow O$ was compared to the alternatives. Thus, probabilities are probabilities of transitions. (In the causally trivial case of a deterministic transition, the probability will be trivially equal to one.) We will bind the notion of probabilities fully to the notion of *causae causantes*: since the framework has room for indeterminism, it should be possible to grade that indeterminism numerically at least in some cases.²

We will reduce all considerations of probabilities of transitions to considerations of probabilities of combinations of basic transitions. This will be done by in some way identifying a transition with its set of *causae causantes*. We call two transitions *causally equivalent* ($I \rightsquigarrow O \stackrel{c}{\equiv} I' \rightsquigarrow O'$) iff they have the same *causae causantes*:

Definition 3 (Causal equivalence of transitions)

$$I \rightsquigarrow O \stackrel{c}{\equiv} I' \rightsquigarrow O' \quad \text{iff} \quad cc(I \rightsquigarrow O) = cc(I' \rightsquigarrow O')$$

The relation $\stackrel{c}{\equiv}$ is an equivalence relation, and an equivalence class may be specified via the set of *causae causantes* of one of its members. In what follows, we will therefore not talk about transitions in general, but about sets of basic transitions.³ The main work that is to be done is to find out how basic transitions combine, and which probabilities should be assigned to such combinations.

Thus, the next section introduces some facts about basic transitions, and the section after that is concerned with probabilities.

2 Basic transitions

2.1 Basics of basic transitions

Basic transitions are transitions from a point event (initial) e_i to one of e_i 's immediate outcomes $H_i \in \Pi_{e_i}$; we write

$$t_i = e_i \rightsquigarrow H_i$$

² It seems to be rather uncontroversial that probabilities are numerically graded modalities; cf. e.g., Popper ([1982], p. 70) and van Fraassen ([1980], p. 198). Once possibilities are there, probabilities are (well, almost) for free.

³ This does not mean that all that is important about a transition is its set of *causae causantes*. For example, epistemically the spatio-temporal location of I and O may play a huge role, and there are also important differences between a basic transition and a non-basic transition that is causally equivalent to it. All we claim in this paper is that on the ontological level, all that is *probabilistically* important about a transition is in its *causae causantes*, taken together.

The different $H_i \in \Pi_{(e_i)}$ partition $H_{(e_i)}$, the set of all histories containing e_i . Points e_i for which $\Pi_{(e_i)}$ has only one member are called *trivial*; at such points, no splitting occurs. A transition $e_i \rightarrow H_{(e_i)}$ with a trivial initial is also called trivial. We will mostly be interested in *non-trivial*, or *indeterministic*, points and transitions. A cause-like locus is always indeterministic by definition, so sets of *causae causantes* consist of non-trivial basic transitions.

We call two basic transitions $t_1 = e_1 \rightarrow H_1$ and $t_2 = e_2 \rightarrow H_2$ *initial-equivalent* iff their initials are the same ($e_1 = e_2$); in this case, we write $t_1 \stackrel{i}{\equiv} t_2$. This is again an equivalence relation, and the equivalence class of t_i , $[t_i]$, is the set of all basic transitions with initial e_i . Since $[t_i]$ is completely specified by the initial e_i , we also write the equivalence classes as $[e_i]$.⁴

$$[t_i] = [e_i] = \{e_i \rightarrow H \mid H \in \Pi_{e_i}\}$$

We call an initial e_i *finitely splitting* iff $[e_i]$ is finite. On the set of basic transitions, we can define a partial ordering relation as follows:

$$t_i < t_j \quad \text{iff} \quad e_i < e_j \quad \text{and} \quad H_{(e_j)} \subseteq H_i.⁵$$

Thus, the transition t_j is above t_i ($t_i < t_j$) if t_j can occur in the outcome H_i of t_i . For this, it is *not* enough that $e_i < e_j$: since at e_i , histories split, e_j must also occur in the outcome H_i of e_i that t_i describes.

The ordering of transitions sheds some light on when two transitions are compatible. This will now be spelled out in terms of consistent sets of transitions.

2.2 Sets of basic transitions

Sets of basic transitions will be the basic building blocks for the probability spaces to be defined later on: they will constitute the elements of the sample space. However, we cannot use just any sets of basic transitions for this purpose. Some sets are such that they cannot occur in one history — they are causally impossible. The most useful sets of basic transitions are the consistent ones, where consistency is something like causal possibility.

2.2.1 Consistency of sets of basic transitions

The notion of consistency is defined as follows:

Definition 4 (Consistent set of basic transitions) *A set $T = \{t_1, t_2, \dots\}$ of basic transitions ($t_i = e_i \rightarrow H_i$) is consistent iff all the outcomes have a history in*

⁴ This notion of initial-equivalence of basic transitions is to be distinguished from the notion of causal equivalence of transitions employed above. Note also that on $[t_i]$, probabilities can be easily added, since we can introduce a natural set-theoretic Boolean algebra structure on Π_{e_i} . This will be our starting point later on.

⁵ Note that given $e_i < e_j$, if $H_{(e_i)} \cap H_i \neq \emptyset$, then $H_{(e_j)} \subseteq H_i$ by transitivity of \equiv_{e_i} .

common, i.e., iff $\bigcap_{t \in T} H_t \neq \emptyset$. (The empty set \emptyset is considered consistent, since the intersection of an empty set of histories, $\bigcap_{t \in \emptyset} H_t$, is by definition taken to be the set \mathfrak{H} of all histories.)

By this definition, any subset of a consistent set of basic transitions is consistent. A good thing to note is that the set of basic transitions that we identified with a general transition $I \rightarrow O$ is consistent in the sense just defined.

Lemma 1 *For any transition $I \rightarrow O$, the set of its causae causantes is consistent.*

Proof: Let $I \rightarrow O$ be given, and let $T = cc(I \rightarrow O)$. By definition, any $e_i \rightarrow H_i = t_i \in T$ satisfies $e_i < O$ and $H_i = \Pi_{(e_i)} \langle O \rangle$, so that $H_{\langle O \rangle} \subseteq H_i$. Thus, $H_{\langle O \rangle} \subseteq \bigcap_{t \in T} H_t$. As $H_{\langle O \rangle} \neq \emptyset$, T is consistent. \square

Before we can move on, we need to state, and assume, an important postulate and a convenient limitation.

2.2.2 Limitations: finiteness and exclusion of modal funny business

The theory to be developed here will be limited in two ways. The first limitation is for the sake of convenience and will have to be removed in a further development of the theory: we will assume that all sets of *causae causantes* and all probability structures that we are going to deal with are finite. This is by no means a ‘natural’ limitation, i.e., we do not assume that there are any deep metaphysical reasons for this limitation. It is just that sticking to the finite case will allow us to focus on the problem of combining causality and probability without having to cope with problems of infinite probability structures all along.⁶

The second limitation is fundamental. In the general theory of branching space-times, there may occur a phenomenon that Nuel Belnap has aptly called ‘modal funny business’: roughly, some combinatorially possible histories may be absent. Since we rely on smooth combinatorics, we will have to exclude models involving modal funny business from our considerations. This seems not to be much of a constraint, however, since a modally ‘funny’ model may be mimicked probabilistically by a model in which the combinatorics are smooth, and formerly missing histories are excluded via probability zero.

Belnap ([2003]) proves the equivalence of four different notions of modal funny business. For our purposes, it is best to use the following notion:

Definition 5 (Modal funny business) *In a model of branching space-times, there is primary space-like related modal funny business iff there are two initial events (nonempty sets of point events such that each set is part of one history)*

⁶ Cf. note 8 below for some comments on infinite structures.

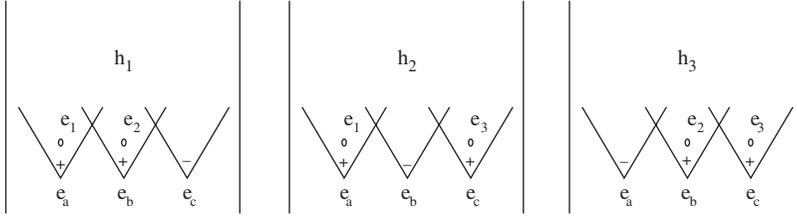


Figure 1. Effects of modal funny business.

I_1, I_2 such that $I_1 \text{ SLR } I_2$ ($\forall e_1 \in I_1 \forall e_2 \in I_2 e_1 \text{ SLR } e_2$), but one primary outcome $\Pi_{I_1}\langle h_1 \rangle$ of I_1 has an empty intersection with one primary outcome $\Pi_{I_2}\langle h_2 \rangle$ of I_2 .⁷

Figure 1 illustrates what may happen if there is modal funny business. The three histories depicted are all the histories there are in the model. Consider first the events $e_a, e_b,$ and e_c that occur in all three histories. Each of these is an indeterministic point event with two outcomes, which are labelled ‘+’ and ‘-’. There is modal funny business according to Definition 5: for example, for $I_1 = \{e_a\}, I_2 = \{e_b\}$, the intersection of the ‘-’ outcomes is empty; among the three histories, none contains both $e_a \rightsquigarrow -$ and $e_b \rightsquigarrow -$. In the case depicted, there are also three pairwise space-like related points $e_1, e_2,$ and e_3 that are nonetheless not all contained in one single history: each pair, e.g., $e_1, e_2,$ is incomparable, and there is a history, e.g., $h_1,$ containing both points, witnessing the space-like relatedness. But contrary to what one might assume, none of the three histories of the model contains all three points e_1, e_2 and e_3 . In the absence of modal funny business, such a situation is ruled out: given no modal funny business according to Definition 5, for any finite set of pairwise SLR points there is a history containing the whole set, and any set of transitions with pairwise SLR initials has a nonempty common outcome.

Lemma 2 *If there is no modal funny business according to Definition 5, the following holds: (a) Let $E = \{e_1, \dots, e_n\}$ be such that for $e_i, e_j \in E,$ if $e_i \neq e_j,$ then $e_i \text{ SLR } e_j.$ Then $H_{[E]} \neq \emptyset,$ i.e., there is a history containing all the points in $E.$ (b) Furthermore, let $T = \{t_1, \dots, t_n\}$ be a set of transitions ($t_i = e_i \rightsquigarrow H_i$) such that the set of initials $E = \{e_i \mid t_i \in T\}$ fulfils the conditions of (a) and such that no two transitions from T have the same initial. Then T is consistent, i.e., the transitions in T have a common outcome: $\cap_{t_i \in T} H_i \neq \emptyset.$*

Proof: We prove only (b), from which one may obtain (a) as follows: if only E is given, select arbitrary outcomes $H_i \in \Pi_{e_i}$ for each $e_i \in E.$ If then $\cap_{t_i \in T} H_i \neq \emptyset,$ we have $H_{[E]} \neq \emptyset$ since $H_i \subseteq H_{(e_i)}.$ So assume that the premises of (b) hold, but that $\cap_{t_i \in T} H_i = \emptyset.$ Let $T' \subseteq T$ be a subset of T that is maximal with respect

⁷ In the infinite case, this definition will have to be altered; cf. note 8 below for some indications of what needs to be changed.

to $\bigcap_{t_i \in T'} H_i \neq \emptyset$. Now let $t_j \in T - T'$. Set $E' = \{e_i \mid t_i \in T'\}$. By assumption, the elements of E are pairwise SLR, so the sets $I_1 = E'$, $I_2 = \{e_j\}$ are initials, and the sets $\bigcap_{t_i \in T'} H_i$ and H_j are primary outcomes, fulfilling the requirement of Definition 5. Now since there is no modal funny business by assumption, that Definition implies $\bigcap_{t_i \in T'} H_i \cap H_j \neq \emptyset$, contradicting the maximality of T' . \square

In what follows, we will assume that there is no modal funny business:

Postulate 1 (No modal funny business) *There is no modal funny business according to Definition 5.*

2.2.3 Characterisation of consistency in the absence of modal funny business

Given finite structures and no modal funny business, the notion of consistency of a set of basic transitions can be expressed in different ways:

Lemma 3 *Let $T = \{t_1, \dots, t_n\}$, $t_i = e_i \rightsquigarrow H_i$. The following conditions are equivalent:*

1. T is consistent.
2. There is a history to which all the e_i belong ($\bigcap_{t_i \in T} H_{(e_i)} \neq \emptyset$). Furthermore,
 - (a) if $e_i = e_j$, then $H_i = H_j$, and (b) if $e_i < e_j$, then $H_{(e_i)} \subseteq H_j$.
3. For any $t_i, t_j \in T$, exactly one of the following holds:
 - (a) $e_i = e_j$ and $H_i = H_j$ (i.e., $t_i = t_j$) or
 - (b) $t_i < t_j$ or
 - (c) $t_j < t_i$ or
 - (d) e_i SLR e_j .

Proof: $1 \Rightarrow 2$: For all t_i , $H_i \subseteq H_{(e_i)}$, and $\bigcap_{t_i \in T} H_i \neq \emptyset$ by (1). So there is a history to which all the e_i belong. If $e_i = e_j$, then $H_i \cap H_j \neq \emptyset$ iff $H_i = H_j$ (the family of outcomes of e_i is a partition of $H_{(e_i)}$). If $e_i < e_j$, then $H_{(e_i)} \cap H_i = \emptyset$ would imply $H_j \cap H_i = \emptyset$ (since $H_j \subseteq H_{(e_j)}$), meaning that T would be inconsistent.

$2 \Rightarrow 3$: Assume T violates (3), i.e., there are $t_i, t_j \in T$ such that none of (3a)–(3d) applies. This is possible in four ways: (a) $e_i = e_j$, but $H_i \neq H_j$ – this violates (2); (b) $e_i < e_j$, but $H_{(e_i)} \cap H_i = \emptyset$, again violating (2); (c) for $e_j < e_i$ symmetrical to (b); (d) $e_i \not\text{SLR } e_j$ and $e_j \not\text{SLR } e_i$ and not e_i SLR e_j – this means that there is no history to which both e_i and e_j belong, again violating (2).

$3 \Rightarrow 1$: Let $T' := \{t_i \in T \mid \neg \exists t_j \in T \ t_i < t_j\}$ be the set of maximal elements of T . It will be sufficient to establish that $\bigcap_{t_i \in T'} H_i \neq \emptyset$, since if $t_i < t_j$, $H_j \subseteq H_i$. Now for $t_i, t_j \in T'$, if $t_i \neq t_j$, then we have e_i SLR e_j ; $e_i = e_j$ is ruled out by (3a), and $t_i < t_j$ or $t_j < t_i$ is impossible since T' is the set of the maximal elements. Thus we have a set of pairwise space-like related points with outcomes

H_i specified, and by Lemma 2 (b), there is a history in the intersection of all outcomes. \square

The following Lemma characterises the special status of the maximal elements: their outcomes may be changed at will without violating consistency.

Lemma 4 *If $T = \{t_1, \dots, t_n\}$ is consistent and $t_0 \in [t_i]$, $t_0 \neq t_i$, then $T' = (T - \{t_i\}) \cup \{t_0\}$ is consistent iff t_i is maximal in T .*

Proof: ‘ \Leftarrow ’: If t_i is maximal, then the set with t_0 substituted is also consistent by Lemma 3 (2).

‘ \Rightarrow ’: Let t_i not be maximal in T , i.e., there is $t_j \in T$ s.t. $t_i \leq t_j$, i.e., $e_i \leq e_j$ and $H_{(e_j)} \cap H_i \neq \emptyset$. Since Π_{e_i} is a partition of $H_{(e_i)}$, $H_{(e_j)} \cap H_0 = \emptyset$. By Lemma 3 (2), the set with t_0 substituted is thus inconsistent. \square

3 Causal probability theory

We have now prepared the ground for the introduction of probabilities into the framework of causation in branching space-times. By the close link between causation and probability that was assumed at the outset, the objects that figure as causes – sets of basic transitions – will be the objects on which probabilities will be defined. In order to have probabilities, we need probability spaces. We rehearse the usual definition:

Definition 6 (Probability space) *A probability space is a triple $\langle A, F, \mu \rangle$, where A is a nonempty set (the ‘sample space’), F is a Boolean σ -algebra (sometimes called ‘event algebra’) of subsets of A (with zero element $\mathbf{0} = \emptyset$ and unit element $\mathbf{1} = A$), and μ is a normalised, countably additive measure on F , i.e., (i) $\mu(\mathbf{1}) = 1$ and (ii) if $(f_i)_{i \in J}$ is an at most countable family of disjoint elements of F , then $\mu(\cup_{i \in J} f_i) = \sum_{i \in J} \mu(f_i)$.*

As stated above, we will consider *finite* probability structures exclusively. Thus, we will assume in the following that all probability spaces that we deal with are finite. In a finite probability space, F is just the usual set-theoretic algebra of subsets of A , with the elements of A playing the role of atoms, or basic alternatives.

The basic questions of this section may be characterised as follows:

Given a transition $I \rightsquigarrow O$, represented by a finite, consistent set T of basic transitions: in which probability space can we speak of a ‘probability for T ’? That is:

1. What is the sample space A that contains T and its causal alternatives?
2. What is the Boolean σ -algebra F ?

3. What is the measure μ ; which general constraints hold for μ ?
4. How do the measures in different probability spaces relate?

We will motivate our answers to these questions through a number of simple examples.

3.1 Some simple cases

The simplest example is given by a trivial transition: its set of *causae causantes* is empty. In this case, we have $A = \{\emptyset\}$, F consists of two elements only (A and the empty set), and $\mu(A) = 1$. The first non-trivial example is given by a singleton set of transitions.

3.1.1 $T = \{t_1\}$

In this case, t_1 is a basic transition with initial e_1 and outcome $H_1 \in \Pi_{e_1}$. The probability space that needs to be considered is obvious: the sample space is naturally taken to be $[t_1]$, i.e., the set of all basic transitions with initial e_1 . Our finiteness assumption requires that $[t_1]$ be finite. This is not guaranteed, and it is certainly not adequate in all situations, but there exist well-known measure-theoretic ways of dealing with infinities in this case (cf. note 8 below).

Given finite $A = [t_1]$, F is just the set of all subsets of A . The measure μ is uniquely characterised by its value on the singleton sets, since any member of F is the disjoint union of a finite number of singletons. This is all standard probability theory, with no hint of causal constraints.

3.1.2 $T = \{t_1, t_2\}, e_1 \text{ SLR } e_2$

The simplest case of a set T with two elements is the case in which the initials are space-like separated. (We do not consider the case that T is inconsistent at the moment, since that cannot happen for T derived from a given transition $I \mapsto O$; cf. Lemma 1.)

In the present case, T combines two transitions. For single transitions, we already know how to construct probability spaces – simply take the set $[t_i]$ to be the sample space, as outlined above. From the two probability spaces $P_1 = \langle A_1, F_1, \mu_1 \rangle$ and $P_2 = \langle A_2, F_2, \mu_2 \rangle$, we may then form the *product space*, $P = \langle A, F, \mu \rangle$, where $A = A_1 \times A_2$ is the Cartesian product of A_1 and A_2 , i.e., the set of ordered pairs $\langle a_1, a_2 \rangle$ with $a_i \in A_i$, $i = 1, 2$, and $F = F_1 \times F_2$ is the respective product Boolean algebra. For the measure μ , we at least require the so-called *marginal property*:

Definition 7 (Marginal property for pairs) *A measure μ on the product Boolean algebra $F = F_1 \times F_2$ has the marginal property iff (i) for all $f_1 \in F_1$, $\mu(\langle f_1, \mathbf{1}_2 \rangle) = \mu_1(f_1)$ and (ii) for all $f_2 \in F_2$, $\mu(\langle \mathbf{1}_1, f_2 \rangle) = \mu_2(f_2)$.*

The marginal property means that if one ignores the outcome of one initial by considering the certain outcome $\mathbf{1}_i = A_i$ (for which $\mu_i(\mathbf{1}_i) = 1$), one gets back the single probability. Requiring that μ satisfy the marginal property does not fix μ there are many ways open. One specific way is taking the *product measure* μ_\times :

$$\mu_\times(\langle f_1, f_2 \rangle) = \mu_1(f_1) \cdot \mu_2(f_2)$$

The other way is to allow for *correlations* between outcomes of e_1 and e_2 by using a different measure.

Here comes a crucial question: given that e_1 and e_2 are space-like related (and thus, causally separated), should we require that $\mu = \mu_\times$? We believe that no, since quantum-mechanical correlation experiments give strong evidence that there can be correlations between outcomes of space-like separated events. In order to remain flexible, our theory should allow both for correlations and for the product measure.

The product construction described above is standard in probability theory. A different perspective on that construction will prove to be illuminating when it comes to generalisations. By Lemma 2, any combination of one outcome of e_1 and one outcome of e_2 is consistent: each $\{t'_1, t'_2\}$ with $t'_1 \in [t_1]$, $t'_2 \in [t_2]$, is consistent. We may base our probability space on such sets and define:

$$A = \{ \{t'_1, t'_2\} \mid t'_1 \in [t_1], t'_2 \in [t_2] \}$$

By the assumption of finitude, F may again be the usual set algebra, and it will suffice to define the measure μ on the atoms, i.e., on the elements of A . Since in this construction, we did not start with ‘single’ probability spaces P_1 and P_2 , we cannot state the marginal property as before. However, we may require some form of marginal property once two ‘single’ probability spaces with respective measures μ_1 and μ_2 are given in addition:

Definition 8 (Marginal property for sets) *Let F be the Boolean algebra of sets of transitions as introduced above. A measure μ on F has the marginal property iff (i) for all $a_1 \in A_1$, $\mu(\{ \{a_1, a_2\} \mid a_2 \in A_2 \}) = \mu_1(\{a_1\})$ and (ii) for all $a_2 \in A_2$, $\mu(\{ \{a_1, a_2\} \mid a_1 \in A_1 \}) = \mu_2(\{a_2\})$.*

This is quite similar to Definition 7: Since

$$\mu(\{ \{a_1, a_2\} \mid a_2 \in A_2 \}) = \sum_{a_2 \in A_2} \mu(\{ \{a_1, a_2\} \}),$$

requiring that that sum be equal to $\mu_1(\{a_1\})$ again means that by ignoring the outcome of e_2 (by summing over all possibilities), one gets back the ‘single’ measure.

In the present case (e_1 SLR e_2), the two perspectives mentioned are provably equivalent, since there is an isomorphism between the two probability spaces mentioned, via

$$\langle a_1, a_2 \rangle \leftrightarrow \{a_1, a_2\}$$

The important question is which of the two perspectives generalises. In classical probability theory, combinations of chance set-ups are adequately described via product constructions. It turns out, however, that once causal notions enter the picture, product constructions are no longer adequate. This is shown by the next example.

3.1.3 $T = \{t_1, t_2\}$, $t_1 < t_2$

Apart from the case e_1 SLR e_2 , which has already been considered, and the case $t_2 > t_1$, which is symmetrical to the present one, this is the only possible case of a consistent set of two transitions (cf. Lemma 3 (3)). As in the previous case, two chance set-ups are combined. The attempt to apply a product construction will result in an instructive failure.

As in the case for e_1 SLR e_2 , one may try to start from two probability spaces $\langle A_i, F_i, \mu_i \rangle$ with $A_i = [t_i]$, $i = 1, 2$. The product space $\langle A, F, \mu \rangle$ is definable as before. However, that probability space is not adequate for the causal situation assumed. Consider a pair $\langle t'_1, t'_2 \rangle$ with $t_1 \in [t_1]$, $t'_1 \neq t_1$, $t'_2 \in [t_2]$. That pair corresponds to a causally impossible 'alternative' to T , since by Lemma 4, there is no history in which both the outcome H'_1 of t'_1 and the initial e_2 of t'_2 occur together. The only reasonable thing to do probabilistically seems to be to require that $\mu(\langle t'_1, t'_2 \rangle) = 0$ in that case (i.e., in case $t'_1 \neq t_1$). But then, given the marginal property, we get

$$\mu_1(t'_1) = \mu(\langle t'_1, \mathbf{1}_2 \rangle) = \sum_{t'_2 \in [t_2]} \mu(\langle t'_1, t'_2 \rangle) = 0 \quad \text{for } t'_1 \neq t_1$$

and by normalisation of μ_1 , we get $\mu_1(t_1) = 1$: e_1 is bound to have outcome H_1 . That is certainly not warranted.

Consider a real life example. Assume that at e_1 you decide at random whether to go to the races or not (possibly by flipping a fair coin): $[e_1] = \{e_1 \rightsquigarrow H_1^1, e_1 \rightsquigarrow H_1^2\}$, $\mu(\text{go}) = \mu_1(e_1 \rightsquigarrow H_1^1) = 0.5$. If you are at the races (e_2 such that $e_1 < e_2$ and $H_{(e_2)} \subseteq H_1^1$), you either win ($e_2 \rightsquigarrow H_2^1$) or lose ($e_2 \rightsquigarrow H_2^2$); suppose $\mu_2(\text{win}) = \mu_2(e_2 \rightsquigarrow H_2^1) = 0.2$. The product space contains the four atoms

$$\langle \text{go}, \text{win} \rangle, \langle \text{go}, \text{lose} \rangle, \langle \text{stay}, \text{win} \rangle, \text{ and } \langle \text{stay}, \text{lose} \rangle$$

The latter two are causally impossible: if you do not go to the races, you can neither win nor lose. By the argument given above, it seems to follow that therefore you go with probability 1. This is absurd.

A way out of this quandary is to use the alternative to the product space construction that was introduced above: consider only the consistent combinations of outcomes of e_1 and e_2 and take as the atoms the maximally consistent sets. Thus, we take the sample space of alternatives to be

$$A = \{\{t_1, t'_2\} \mid t'_2 \in [t_2]\} \cup \{\{t'_1\} \mid t'_1 \in [t_1], t'_1 \neq t_1\}$$

Given finiteness, the Boolean algebra F is again the usual set-theoretic algebra, and μ is uniquely specified by its value on the atoms. Each element of A is a maximally detailed description of a consistent outcome involving the initials e_1 and (possibly) e_2 , and each element of F is a set of such consistent outcomes. A then contains exactly the causal alternatives to T . In terms of the races example, we consider as atomic outcomes only the three (instead of four) maximally consistent sets (not pairs) of transitions,

$$\{\text{go, win}\}, \{\text{go, lose}\}, \{\text{stay}\}$$

Note that while the set $\{\text{go}\}$ is consistent, it is not maximally so and is thus not considered to be an atomic possibility. We may, however, identify that set with an element of F : $\{\text{go}\}$ says the same as $\{\{\text{go, win}\}, \{\text{go, lose}\}\}$ if you go, you either go and win, or you go and lose. A kind of marginal property is thus already built into our framework, since the probability of $\{\text{go}\}$ evaluates as

$$\mu(\{\text{go}\}) = \mu(\{\{\text{go, win}\}, \{\text{go, lose}\}\}) = \mu(\{\{\text{go, win}\}\}) + \mu(\{\{\text{go, lose}\}\})$$

by probability theory alone. We will spell this out by defining the notion of a representative below.

The framework sketched for our example easily generalises to arbitrary finite causal probability structures, as will be shown in the following.

3.2 General causal probabilities

We start with a given consistent set T of transitions, e.g., the set of *causae causantes* of a given transition $I \mapsto O$. Since we only consider finite structures in this paper, we require that (i) T be finite and that (ii) for each $t_i \in T$, $[t_i]$ be finite.⁸

⁸ It turns out that while it is rather simple to generalise with respect to requirement (ii), using standard tools of measure theory (e.g., Borel sets), it is much harder to generalise with respect to (i). In fact, infinite T creates difficulties both for considerations of probabilities and for the notion of 'modal funny business'. With respect to the latter, it turns out that Definition 5 is no longer adequate for guaranteeing 'smooth combinatorics'. In order to exclude combinatorially weird cases, it seems best to use the step from condition (3) to condition (1) in Lemma 3 as the mark of 'no funny business'. The whole area merits further study. With respect to probabilities,

The causal probability space $P_T = \langle A_T, F_T, \mu_T \rangle$ that is adequate to describe causal alternatives to T is constructed as follows:

1. Let $\tilde{T} = \cup_{t_i \in T} [t_i]$, the set of all alternatives to the elements of T .
2. Let $A_T = \{T' | T' \subseteq \tilde{T}, T' \text{ maximally consistent}\}$.
3. Let F_T be the set-theoretic Boolean algebra over A_T .
4. The measure μ_T may be any normalised measure on F_T .
5. The sought-for probability space is $P_T = \langle A_T, F_T, \mu_T \rangle$.

As one expects, the set T of transitions turns out to be an element of A_T since T is consistent, and by Lemma 3 (3a), no other element of \tilde{T} can be consistently added to T . The given construction does not in any way single out T from among the other elements of A_T . This is as it should be, since a probability space does not contain a designated element to stand for ‘the real outcome’. Thus, the probability space P_T can be constructed once the initials of all transitions in T are known: it suffices to start from a given consistent set of point events E ; the corresponding probability space may then be denoted $P_E = \langle A_E, F_E, \mu_E \rangle$. Since this way of specifying causal probability structures can be generalised most easily, we will mainly be talking about probability spaces P_E with a given set E of initials. It is even possible to build probability spaces from given *inconsistent* sets of point events; we will consider this generalisation below.

In standard probability theory, it is customary to require that different measures in different probability spaces respect the marginal property, which was formulated for pairs as Definition 7 above. Since causal probability structures are not product spaces, we need some way of expressing the fact that we wish to consider the probability of ‘the same thing’ in different probability spaces.

3.2.1 Representation in different probability spaces

Reconsider the races example from Section 3.1.3. Going was in some sense an alternative, but it was not a maximally specific alternative, and thus not an element of A_T . Still, there was a place for it in F_T , viz., as the set $\{\{\text{go, win}\}, \{\text{go, lose}\}\}$. More generally, we may ask: given a consistent set of basic transitions S , i.e., an element of the sample space A_E of some causal probability space P_E , and a causal probability space $P_{E'}$, how is S represented in $P_{E'}$?

if T is infinite, i.e. involves infinitely many initials, then there may be an infinite SLR set of initials, and there may be an infinite chain of transitions in T . In the first case, the tools of standard probability theory for infinite product spaces (cylinder sets, zero-one laws) will apply. In the second case, the situation appears to be more challenging. Still, that case is very interesting in view of the fact that, e.g., in the modal theory of agency, ‘busy choice sequences’ are analysed whose probabilistic equivalent exactly requires a probability theory for infinite chains of transitions. That question, too, certainly merits further study.

That question may be split in two: (i) When is S representable in $P_{E'}$? (ii) If S is representable, what is the representative? The answer to the first question is straightforward: the initials of S need to be in E' . Thus we define:

Definition 9 (Representability of S in $P_{E'}$) A consistent set $S \in A_E$ of basic transitions is representable in $P_{E'}$ iff $E \subseteq E'$.

When S is representable, we have seen that we should not expect $S \in A_{E'}$, since $A_{E'}$ contains maximally consistent sets relative to E' only, and S may not be maximal. Nor will we have $S \in F_{E'}$, since $F_{E'}$ is a set of (maximally consistent) sets of transitions. However, the following definition captures the required notion:

Definition 10 (Representative of S in $P_{E'}$) The representative of S in $P_{E'}$, $S_{E'}$, is defined to be the set of all maximally consistent sets extending S :

$$S_{E'} = \{T' \in A_{E'} \mid S \subseteq T'\}$$

Thus we have $S_{E'} \in F_{E'}$. In order to visualise what is going on, it may be good to divide the definition into two steps, one downward and one upward, as follows:

1. $S' = S \cup \{e' \rightarrow \Pi_{e'}\langle e \rangle \mid e \rightarrow H \text{ minimal in } S, e' < e, e' \in E'\}$

In words: S' is derived from S by ‘downward closure in E' ’, i.e., by adding all transitions below S that are consistent with S and that have initials in E' .

2. $S_{E'} = \{T' \in A_{E'} \mid S' \subseteq T'\}$

In words: the representative of S in $P_{E'}$ is the set of all maximally consistent sets extending S' . Since S' is downward closed in E' , this step will only add elements on top of S .

The concept of a representative allows us to state, and prove, that the probability of a set of transitions is always smaller than, or equal to, the probability of its subsets:

Lemma 5 Let S be representable in P_E , and let $S' \subseteq S$. Then $\mu_E(S_E) \leq \mu_E(S'_E)$.

Proof: As $S' \subseteq S$, S' is representable in P_E as well. From the definition of S_E , one reads off $S_E \subseteq S'_E$. The claim then follows from the additivity of μ_E . \square

Representing elements of the event algebra So far, we have only considered the representation of consistent sets of transitions, i.e., elements of the sample space of some probability space. More generally, we may be interested in defining a representative for a set of consistent sets of transitions, e.g.,

an element of the event algebra of some probability space. The respective definitions are natural generalisations of the two definitions from above:

Definition 11 (Representability) *A set $f \in F_E$ of sets of basic transitions is representable in $P_{E'}$ iff $E \subseteq E'$.*

Thus, f is representable in $P_{E'}$ iff each $S \in f$ is representable in $P_{E'}$.

Definition 12 (Representative of f in $P_{E'}$) *The representative of f in $P_{E'}$, $f_{E'}$, is defined to be the union of all representatives of elements of f :*

$$f_{E'} = \cup_{S \in f} S_{E'}$$

By this definition, we have $f_{E'} \in F_{E'}$: an element of the event algebra of P_E is represented by an element of the event algebra of $P_{E'}$.

3.2.2 Adapting the marginal property

As in standard probability theory, we want to say that the probability of ‘something’ is the same in all probability spaces representing the something. Taking the ‘something’ to be a consistent set of basic transitions, one might think that we would like to have the following alternative marginal property:

(BAD-MP) If $S \in A_E$ is representable in $P_{E'}$, then $\mu_E(S) = \mu_{E'}(S_{E'})$.

In standard probability theory, this is a good way to describe the marginal property. In our framework, it will not do. Consider $S = \{e_1 \mapsto H_1\}$ and $E = \{e_1\}$, $E' = \{e_0, e_1\}$ with $e_0 < e_1$. In this situation, the probability of S will normally be greater than the probability of $S_{E'} = \{\{e_0 \mapsto \prod_{e_0} \langle e_1 \rangle, e_1 \mapsto H_1\}\}$, since in $P_{E'}$, first e_0 has to have the right outcome and then e_1 has to have the right outcome, whereas in P_E , only e_1 has to have the right outcome. Standard probability theory with the product construction does not leave room for picturing more than one ‘layer’ of initials, but the framework that is presented here does. In order to arrive at a correct formulation of the marginal property, we need to require that $P_{E'}$ be an *extension* of P_E , which notion is defined as follows:

Definition 13 (Extension) *A causal probability space $P_{E'}$ is an extension of P_E iff (i) $E \subseteq E'$ and (ii) there is no $e' \in E' - E$ and $e \in E$ such that $e' < e$.*

In other words, $P_{E'}$ is an extension of P_E iff E' is a superset of E and all the new elements of E' are properly maximal with respect to E . If $P_{E'}$ extends P_E , the problematic situation described above cannot occur. Thus, the following is an adequate formulation of the marginal property for sets of basic transitions in our causal framework:

(MP) If $S \in A_E$ is representable in $P_{E'}$

and if $P_{E'}$ extends P_E , then $\mu_E(S) = \mu_{E'}(S_{E'})$.

We postulate that the marginal property should hold in Our world:

Postulate 2 *The marginal property (MP) holds for all probability assignments.*

Postulate 2 is already sufficient to ensure that an analogue of (MP) holds for elements of event algebras as well. We state this result as a Lemma:

Lemma 6 *Let $f \in F_E$, and let $P_{E'}$ extend P_E . Given (MP), it follows that $\mu_E(f) = \mu_{E'}(f_{E'})$.*

Proof:

$$\mu_E(f) = \sum_{S \in f} \mu_E(S) = \sum_{S \in f} \mu_{E'}(S_{E'}) = \mu_{E'}(\cup_{S \in f} S_{E'}) = \mu_{E'}(f_{E'}). \quad \square$$

3.2.3 Markov property

The motivation for the marginal property was purely *structural*: ‘the same thing’ has to have the same probability, however represented. The causal structure only entered in spelling out what ‘representing the same thing’ might mean. We now consider a property of probability spaces that is of *causal* origin and that accordingly does not have a direct counterpart in standard probability theory. We rely on the intuition that the probability of a transition is something like the probability of the outcome, given the initial. We will suggest that in some specific cases, this may be read as a conditional probability.

We first introduce the notion of *layering*:

Definition 14 *Let T be consistent. We say that T is layered iff there is a proper subset T' of T such that for all $t' \in T'$ and for all $t \in T - T'$, $t < t'$. In this case, we say that T' forms an upper layer of T .*

Not all T are layered. Consider Figure 2A, where $e_2 < e_3$, but e_1 SLR e_2 and e_1 SLR e_3 . If T is layered, as in Figure 2B and C, an upper layer may be uniquely determined, as in (B), or it may not be unique, as in (C).

Transitions in an upper layer of layered T all are located above all the other transitions. Thus there can be no causal influence from the transitions in $T - T'$ on one of the outcomes of T' that was not already present when T' occurs. Given that T' occurs, the occurrence of $T - T'$ has been accounted for completely. This observation motivates the following postulate:

Postulate 3 (Markov property) *Let $T = R \cup S$ be consistent and layered, and let R be an upper layer of T . Then all measures μ_T on P_T satisfy*

$$\mu_T(T) = \mu_R(R) \cdot \mu_S(S)$$

If $\mu_S(S) \neq 0$, this may be phrased in terms of conditional probabilities:

$$\mu_R(R) = \mu_T(T \mid S) = \mu_T(\{T\} \cap S_T) / \mu_T(S_T) = \mu_T(T) / \mu_S(S),$$

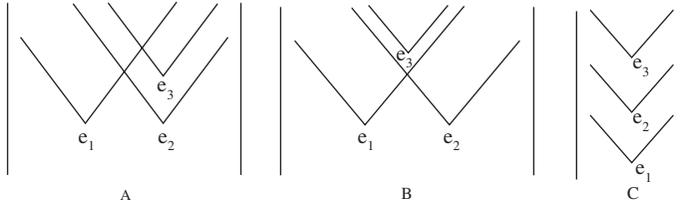


Figure 2. A: not layered; B and C: layered.

where we have employed $T \in \mathcal{S}_T$ (since $S \subseteq T$) and $\mu_S(S) = \mu_T(S_T)$ (by the marginal property).

Note that we are here dealing with an *ontologically motivated* Markov condition that applies to single cases. In connection with graphical models and Bayesian Networks, one assumes a so-called ‘causal Markov condition’ that applies to ‘variables’. We will give some brief remarks on the interrelation between the framework presented here and the framework of Bayesian Networks at the end of Section 3.2.6 below.

3.2.4 Causal probability spaces from inconsistent initials

Technically, it is unproblematic to extend the probability space construction described at the beginning of Section 3.2 to any set E of initials:

1. Form the set $\tilde{T} = \cup_{e \in E} [e]$.
2. Use the set of all maximally consistent subsets T' of \tilde{T} as sample space A_E .
3. Form the event algebra F_E ; in the finite case that is assumed here for simplicity, use the set-theoretic algebra of subsets of A_E .
4. Define some measure μ_E on F_E , e.g., by assigning values to the atoms $T' \in A_E$ such that $\sum_{T' \in A_E} \mu_E(T') = 1$.
5. Set the probability space P_E to be $\langle A_E, F_E, \mu_E \rangle$.

For this construction to be well-defined it is *not* necessary that E be consistent. For example, in the case of $E = \{e_1, e_2\}$ with incompatible e_1 and e_2 , assuming two outcomes $+_1, -_1$ and $+_2, -_2$ for simplicity, A_E contains only singletons:

$$A_E = \{\{e_1 \rightsquigarrow +_1\}, \{e_1 \rightsquigarrow -_1\}, \{e_2 \rightsquigarrow +_2\}, \{e_2 \rightsquigarrow -_2\}\}$$

The crucial question is: under which circumstances can we give a coherent causal reading to the space P_E ? The guiding idea in constructing probability spaces was that A_E should describe alternatives. In the example, this reading appears questionable: in what sense might the occurrence of $e_1 \rightsquigarrow +_1$ be an alternative to the occurrence of $e_2 \rightsquigarrow +_2$? A minimal sense of ‘alternative’ might be defended: They cannot *both* occur, so they are alternatives. But by

looking for a causal reading of probability spaces, we also want a *causal* sense of ‘alternative’, and that sense is not forthcoming in the example: no connection between the alternatives is present. Things would be different if E contained a third element e such that $e < e_1$, $e < e_2$, for then e would provide the sought-for causal link, shared by all the four alternatives.

In the general case, we want to say that A_E is a set of causal alternatives if there is such a causal link in the form of initials that are shared by all the alternatives. Since A_E consists of maximally consistent sets of transitions, the minima of E will be common initials of all alternatives in A_E exactly if they are consistent. We thus give the following definition of when a probability space is causally interpretable as a space of alternatives:⁹

Definition 15 (Causal probability space) *Let E be a finite set of finitely splitting point events. The probability space P_E , defined as above, is a causal probability space iff the minimal elements of E are consistent.*

Given a causal probability space P_E , it will do no harm to add proper maxima to E – the resulting probability space will again be causal, since its minimal elements stay the same. We call the procedure of adding proper maxima ‘fine graining’. For this idea, refer back to the races example from Section 3.1.3. There, $E = \{e_1, e_2\}$, and there was one minimal initial, e_1 (deciding what to do), with outcomes go/stay. In the outcome ‘go’, there was a second initial, e_2 (betting), with outcomes win/lose. By our definition, the corresponding probability space P_E is causal. That feature will be preserved if we give a more fine-grained description of the ‘stay’ outcome of e_1 by adding a new initial e_3 in the ‘stay’ outcome above e_1 . For example, e_3 might have outcomes read/cook. With $E' = \{e_1, e_2, e_3\}$, $P_{E'}$ clearly describes causal alternatives: you either stay and read, or you stay and cook, or you go and win, or you go and lose. Given one alternative, you can tell a causal story of how another alternative might have occurred. Note that this $P_{E'}$ cannot be derived from any consistent set of transitions directly, but only via fine graining. Our formal definition runs as follows:

Definition 16 (Fine graining) *Let E and E' be finite sets of finitely splitting point events. The probability space $P_{E'}$ is a fine-grained version of P_E iff (i) $E \subseteq E'$, (ii) for all $e' \in E' - E$ there is $e \in E$ such that $e < e'$, and (iii) for all $S \in A_E$, the measure is preserved: $\mu_E(S) = \mu_{E'}(S_{E'})$.*

Thus, fine graining means adding new initials that are not new minima, while preserving the measure as far as possible. We can now formulate a

⁹ The definition is limited to the finite case in view of the general finiteness requirement in this paper.

Lemma that connects the notion of a causal probability space with the notion of fine graining.

Lemma 7 *Let E be a set of point events. The following conditions are equivalent:*

1. P_E is a causal probability space (i.e., the minima of E are consistent).
2. The minima of E are pairwise space-like related.
3. P_E is a fine-grained version of a causal probability space P_T , where T is a consistent set of transitions.

Proof: 1 \Rightarrow 2: Since the set of minima is consistent, there is a history h containing them all. Two minima e_1 and e_2 , $e_1 \neq e_2$, cannot be comparable (for then one would not be minimal), so (as witnessed by h) they are space-like related.

2 \Rightarrow 3: Let M be the set of minima of E . For each $e_i \in M$, select an outcome $H_i \in \Pi_{e_i}$, and let $T = \{e_i \mapsto H_i\}$. By Lemma 3 (3d), T is consistent. Then P_E is derived by fine graining from P_T : (i) $M \subseteq E$ and (ii) $E - M$ contains just the non-minimal elements of E . The measure can be adjusted as necessary.

3 \Rightarrow 1: Let P_E be derived from P_T by fine graining, and let M be the set of minima from among the initials of T . By consistency of T , M is consistent, and by the definition of fine graining, M is also the set of minima of E . \square

3.2.5 Random variables and correlations

In standard probability theory, a *random variable* X is a function defined on the sample space A of a probability space P . Usually, but not necessarily, the values of X will be real numbers. For example, with the sample space A being the outcomes of a roll of a die, X might be equal to one for the odd outcomes and equal to zero for the even ones. The probability of a value x of the random variable X is written $pr(X = x)$. For a fair die, in the example case we would thus have $pr(X = 1) = 1/2$.

Random variables may map more than one element of the sample space to the same value, as in the example. The most fine-grained random variables preserve the structure of the sample space via an isomorphism, e.g., $X(a) = a$. For such random variables, we may identify $pr(X = x)$ with $\mu(x)$. This identification is unproblematic in the finite case considered here. Random variables are required to give a definition of correlation and (probabilistic) independence:

Definition 17 (Independence and correlations) *A family X_1, \dots, X_n of random variables defined on a probability space P is called independent iff for any n -tuple $\langle x_1, \dots, x_n \rangle$ of respective outcomes,*

$$pr(X_1 = x_1 \text{ and } \dots \text{ and } X_n = x_n) = pr(X_1 = x_1) \cdot \dots \cdot pr(X_n = x_n)$$

If the family X_1, \dots, X_n is not independent, it is called correlated. Thus, a correlation is of the form

$$pr(X_1 = x_1 \text{ and } \dots \text{ and } X_n = x_n) \neq pr(X_1 = x_1) \cdot \dots \cdot pr(X_n = x_n)$$

Note that pairwise independence does not guarantee independence. It is straightforward to transfer these notions to causal probability structures: a random variable X_E is again taken to be a function defined on the sample space A_E of a probability space P_E . *Basic random variables* are defined on the sample space $A_{\{t\}}$ generated by a single basic transition $t = e \mapsto H$.

As an example, consider rolling a fair die again, as described by a probability space P_E (where $E = \{e\}$ has just one element). The random variable X will be as defined above. Let Y be defined to have the value one for outcome 2, and value zero otherwise. It turns out that X and Y are correlated, since, e.g.,

$$pr(X = 1 \text{ and } Y = 1) = 0 \neq pr(X = 1) \cdot pr(Y = 1) = (1/2) \cdot (1/6)$$

Such trivial correlations must be distinguished from empirically challenging correlations, for which the notion of screening-off is important. Before we move to a discussion of that notion, we give a definition of *independence* and *correlations* for causal probability spaces, which is based on using the most fine-grained random variables:

Definition 18 Let P_E be a finite causal probability space. P_E is called (probabilistically) independent iff for any $T = \{t_1, \dots, t_n\} \in A_E$,

$$\mu_E(T) = \prod_{t_i \in T} \mu_{\{t_i\}}(t_i)$$

P_E is called correlated iff it is not independent.

Thus, we have at our hands two different concepts of correlations: one is about random variables defined on P_E , while the other is about P_E itself. The die example already shows that on an uncorrelated P_E , one can define correlated random variables. In many empirically important cases, such correlations among random variables may be explained by uncorrelated P_E , invoking the notion of (probabilistic) screening-off.

3.2.6 Screening-off

One of the most important links between causality and probability is via the concept of screening-off. Roughly, screening-off means that if there are correlations among causally unconnected variables, then in the common past of these variables there is an event that screens off the correlations, i.e., conditional on which there are no correlations any more. Consider the notorious

barometer example: a falling barometer is correlated with rain, even though neither does the rain cause the barometer to fall, nor does the falling barometer cause the rain. In this situation, we can point to a temporally prior common cause of the correlations: low atmospheric pressure causes both the rain and the barometer's falling; conditional on low atmospheric pressure, the two variables are uncorrelated.

This story is good as far as it goes, and it captures a methodologically important point that may drive empirical research: explaining correlations through the common past.¹⁰ Since we wish to give a reading of screening-off in the framework of this paper, we need to be more explicit with respect to a number of key notions:

1. What are the variables that may be correlated?
2. For which correlations do we expect screening-off?
3. Which notion of location for variables is appropriate?
4. What does the screening-off principle look like? How can the requirement of causal priority and the concept of conditioning on a common cause be expressed?

(1) It seems natural to take 'variables' to be random variables defined on a causal probability space P_E . Thus, we will be dealing with correlations in the sense of Definition 17.

(2) As a famous slogan would have it, correlations cry out for explanation. That may be so, but certainly not all correlations cry for explanations of the screening-off variety. That is already obvious from the die example, where the random variables were just defined so as to yield correlations. Furthermore, any sort of causal influence of one variable on another will make correlations relatively uninteresting, too. For example, in the races example from Section 3.1.3, going is correlated with winning, but this is just so because in order to win, one has to go and bet in the first place. Thus, interesting correlations are among 'causally separated' variables, which must somehow be expressed via the concept of space-like separation, bringing us to the next point.

(3) Random variables do not have locations – they are just functions. Still, for a random variable X defined on P_E , some points $e \in E$ may be important for determining the value of X , while others may be unimportant. The maximal elements of E play a special role here, since the outcomes of the maximal elements determine an element of the sample space (the domain of X) uniquely. Letting E_M be the set of maximal elements of E , we thus define

¹⁰ In fact, the first overt application of a screening-off principle, by Reichenbach ([1956]), was in trying to define the direction of time.

the set V_X of value-determining points for X to be

$$V_X = \{e \in E_M \mid \exists t, t' \in [e] \exists S \in A_E (t \in S \text{ and } X(S) \neq X((S - \{t\}) \cup \{t'\}))\}$$

In words, a maximal point e of E is value-determining for X iff altering the outcome of e can make a difference as to the value of X .

We may now define when two random variables defined in P_E are *causally separated*:

Definition 19 (Causal separation of random variables) *The random variables X and Y defined in P_E are causally separated iff their respective value-determining points are either space-like separated or incompatible, i.e., iff these points are incomparable by \leq :*

$$\forall e_X \in V_X \forall e_Y \in V_Y (e_X \not\leq e_Y \text{ and } e_Y \not\leq e_X)$$

Thus, we will be looking for screening-off for causally separated variables only. (This excludes both the die case and the races example.)

(4) Let X and Y be two causally separated, correlated random variables defined in a causal probability space P_E . If P_E itself is correlated (according to Definition 18), then correlations among variables are to be expected anyhow. However, in the absence of correlations in P_E , a screening-off principle holds. We state this as a Theorem:

Theorem 1 (Screening-off principle)

Let X and Y be two causally separated, correlated random variables defined in a causal probability space P_E , and let $E' = E - E_M$ be the set of non-maximal elements in E . Then either (i) P_E is correlated or (ii) P_E is uncorrelated, and any $S' \in A_{E'}$ screens off X from Y : for any $S' \in A_{E'}$, any $x \in \text{ran}(X)$ and any $y \in \text{ran}(Y)$,

$$\text{pr}(X = x \text{ and } Y = y \mid S') = \text{pr}(X = x \mid S') \cdot \text{pr}(Y = y \mid S')$$

Proof: We need to consider (ii) only. For simplicity's sake we assume that $V_X \cup V_Y = E_M$, i.e., all maximal points in E are relevant for X or for Y . (Otherwise, the unnecessary points may be deleted.) We write $S_x(S')$ for the set $\{S \in A_E \mid S' \subseteq S, X(S) = x\}$, S_y analogously, and S_{xy} for $\{S \in A_E \mid S' \subseteq S, X(S) = x \text{ and } Y(S) = y\}$. Thus, e.g., $S_x(S')$ pools together all ways of continuing S' such as to yield value x for X . Then we have

$$\begin{aligned} \text{pr}(X = x \text{ and } Y = y \mid S') &= \mu_E(\{S \in S_{xy}(S')\}) / \mu_{E'}(S') \\ &= \sum_{S \in S_{xy}(S')} \mu_E(S) / \mu_{E'}(S') \\ &= \sum_{S \in S_{xy}(S')} \prod_{t \in S} \mu_{\{t\}}(t) \end{aligned}$$

The fraction may be replaced by the product in the last step, since P_E is uncorrelated, and thus both $\mu_E(S)$ and $\mu_{E'}(S')$ factor: e.g., $\mu_E(S) = \prod_{t \in S} \mu_{\{t\}}(t)$. We now write $T_x(S')$ for the set of sets of transitions T with initials from V_X for which there is $S \in A_E$, $S' \subseteq S$, $T \subseteq S$, with $X(S) = x$. Thus, $T_x(S')$ pools together all possible ways for the variable X to have value x , given that S' occurred. (The outcomes of $e \in V_Y$ do not change the value of X .) We thus get

$$\begin{aligned} pr(X = x \mid S') &= \mu_E(\{S \in S_x(S')\}) / \mu_{E'}(S') \\ &= \sum_{S \in S_x(S')} \mu_E(S) / \mu_{E'}(S') \\ &= \sum_{S \in S_x(S')} \prod_{t \in S'} \mu_{\{t\}}(t) \\ &= \sum_{T \in T_x(S')} \prod_{t \in T} \mu_{\{t\}}(t) \end{aligned}$$

The last step is licensed since for $e \notin V_X$, all outcomes of e will lead to the same value of X by definition, so that the sum is over all outcomes of e , yielding probability one and thus cancelling e from the sum.

We finally have

$$\begin{aligned} pr(X = x \mid S') \cdot pr(Y = y \mid S') &= \\ \left(\sum_{T \in T_x(S')} \prod_{t \in T} \mu_{\{t\}}(t) \right) \cdot \left(\sum_{T' \in T_y(S')} \prod_{t' \in T'} \mu_{\{t'\}}(t') \right) &= \\ \sum_{T \in T_x(S')} \sum_{T' \in T_y(S')} \prod_{t \in T} \mu_{\{t\}}(t) \cdot \prod_{t' \in T'} \mu_{\{t'\}}(t') &= \\ \sum_{S \in S_{xy}(S')} \prod_{t \in S'} \mu_{\{t\}}(t) &= pr(X = x \text{ and } Y = y \mid S') \end{aligned}$$

In the last step, we have employed the assumption that V_X and V_Y together make up the set of maxima of E , so that any continuation of S' in A_E has 'new' initials either from V_X or from V_Y . \square

Thus, our framework allows us to capture in a mathematically precise way a notion of 'prior screening-off' for correlated random variables: if variables X and Y are causally separated and the causal probability space P_E on which X and Y are defined is uncorrelated, then the non-maximal elements of E collectively act as screeners-off. Even though random variables do not have a space-time location, the Theorem captures the notion of *prior* screening-off, since the values of X and Y are determined by the maximal elements of E , which are causally later than the screeners-off.

Reichenbach seems to have thought that screening-off was a universal principle. Our Theorem, however, allows for a failure of screening-off, viz.,

if P_E itself is correlated. Correlations in P_E may in turn have two very different reasons: (i) It may be that P_E is inadequate to describe the situation at hand, since it contains correlations where a more careful analysis, extending P_E , would find none. P_E may thus be thought of as a ‘surface model’ for the phenomena. In an extended model (made up from copies of P_E after identifying prior screeners-off), there will be no correlations. (ii) It may be that P_E is the right model, and still there are correlations.

A universal principle of screening-off will not rule out possibility (i). In fact, the principle of screening-off, read as a piece of scientific methodology, urges us to replace situations of type (i) with a more satisfactory, uncorrelated extended model. For example, in the barometer case an initial model may include only a ‘barometer’ event and a ‘rain’ event, and the corresponding causal probability space *will* be correlated. It is only through the introduction of the prior ‘pressure’ event that an uncorrelated space can be constructed. However, universal screening-off speaks against possibility (ii). If universal screening-off holds, then a correlated model can never be the right one. It is an empirical question whether a situation of type (ii) actually occurs. Here, the Bell inequalities together with empirical results of quantum correlation experiments strongly suggest a positive answer: in nature, there are correlations among causally separated variables for which there is no screener-off.¹¹ In the light of the strangeness of these correlations, it seems appropriate to call correlations in P_E ‘probabilistic funny business’. Thus, we define:

Definition 20 *A causal probability space P_E exhibits probabilistic funny business iff it is correlated according to Definition 18.*

A short comment on the theory of Bayesian Networks may be in order. While this paper is about deducing probabilistic consequences from causal relations, it is often important to deduce causal relations from statistical data. Since the 1980s, a number of methods have been developed for elucidating that direction of the problem. The framework of Bayesian Networks (Pearl [2000]; Jensen [2001]) has been especially prominent in this respect, and many important applications have been based on that framework. The present framework may be seen as giving a spatio-temporally and mathematically precise reading of the ‘causal variables’ of the Bayesian Network theory in terms of random variables defined on causal probability spaces. It seems plausible to assume that the further development of the present framework will shed light on fundamental questions in the theory of Bayesian Networks,

¹¹ The Bell inequalities were derived by John Bell after Reichenbach’s death; cf. Bell ([1987]) for the original papers dating from 1964 onward. Decisive experiments have been conducted since the 1980s; cf., e.g., Aspect et al. ([1982]). A more thorough analysis of the quantum mechanical case will have to be deferred to a future paper.

such as the question of the status of the ‘causal Markov condition’ assumed there.¹²

In the following section, we will put many of our notions to work.

3.3 Application: probability of suprema of a chain

As an application of our framework, we consider the following problem:

Let I be an initial chain and $Sup := \{\sup_h(I) \mid h \in H_{[I]}\}$ the family of its suprema. This is a kind of chance set-up that Belnap ([1992]) has called ‘indeterminism without choice’: the transition from I to its suprema is generally indeterministic, but no choice in I determines the outcome. With respect to this set-up, two questions arise:

1. In which probability space may we describe all the transitions $I \rightsquigarrow p$, $p \in Sup$, as causal alternatives?
2. What will be the sum of the probabilities of all these alternative transitions? Will our framework yield the expected answer?

$$\sum_{p \in Sup} \mu(I \rightsquigarrow p) = 1$$

It turns out that we need a further Postulate to deal with this problem. Appendix A shows that Postulate 4, due to Weiner ([1997]), is not a consequence of the other Postulates of branching space-times.

Postulate 4 (Weiner) *Suprema retain their order across histories: If I and J are two initial chains and h_1, h_2 two histories both containing both I and J , then (i) $\sup_{h_1}(I) = \sup_{h_1}(J)$ iff $\sup_{h_2}(I) = \sup_{h_2}(J)$; (ii) $\sup_{h_1}(I) < \sup_{h_1}(J)$ iff $\sup_{h_2}(I) < \sup_{h_2}(J)$, and (iii) $\sup_{h_1}(I) > \sup_{h_1}(J)$ iff $\sup_{h_2}(I) > \sup_{h_2}(J)$.*

Given this Postulate, the answers to the two questions are: (1) There is a natural choice for the sought-for probability space. (2) The probabilities in that space do sum to one. We formulate this as a Theorem. At present, since the considerations of probability spaces were limited to the finite case, we will have to assume finiteness here, too.

Theorem 2

Let I be an initial chain, and let $Sup := \{\sup_h(I) \mid h \in H_{[I]}\}$ be the family of its suprema. Let $CC = \cup_{p \in Sup} cc(I \rightsquigarrow p)$. Then, assuming that CC is non-empty, finite and with finitely splitting initials,

1. P_{CC} is a causal probability space,

¹² A further interesting link between the two frameworks is given by the problem of applying the Bayesian Networks approach to single cases; cf. Pearl ([2000], chap. 10) on the problem of ‘the actual cause’.

2. $ACC = \{cc(I \rightarrow p) \mid p \in Sup\}$, and
3. for any measure μ_{CC} on F_{CC} , $\sum_{p \in Sup} \mu_{CC}(I \rightarrow p) = 1$.

Proof: Let I and Sup be given, and let $PCL = \cup_{p \in Sup} pcl(I \rightarrow p)$, $CC = \cup_{p \in Sup} cc(I \rightarrow p)$. We prove the three claims in turn, spending most of our labour on (2). First we need to prove a Lemma:

Lemma 8 *Let e_i be minimal in PCL , and let $p_j \in Sup$. Then $e_i < p_j$.*

Proof: As $e_i \in PCL$, there is $p_i \in Sup$ s.t. $e_i < p_i$. Assume that $e_i \not< p_j$. In that case, it follows from Postulate 4 that no history can contain both e_i and p_j (use $\{e_i\}$ and I as initials and two histories both containing e_i , one containing p_i , the other p_j). Thus, on the assumption that $e_i \not< p_j$, for a history h_i containing e_i and a history h_j containing p_j , the prior choice principle gives a splitting point $e \in h_i \cap h_j$ s.t. $e < e_i$. But then $e < p_i$, and as $p_j \in h_j$, we have $h_j \in H_{[I]}$, and $h_j \perp_e H_{\langle p_i \rangle}$ by transitivity of undividedness. Thus, $e \in pcl(I \rightarrow p_i)$, contradicting the minimality of e_i . \square

We now proceed to prove the Theorem.

(1) If PCL has only one minimum, there is nothing to show. So let e_1 and e_2 , $e_1 \neq e_2$, be minima of PCL . Let $p \in Sup$. By Lemma 8, $e_1 < p$ and $e_2 < p$, so there is a history containing both. Since e_1 and e_2 , being minima, must be incomparable, it follows that they are space-like related. The claim then follows by Lemma 7 (2). \square

(2) ' \subseteq ': For each transition $I \rightarrow p$ with $p \in Sup$, $cc(I \rightarrow p)$ is a maximally consistent subset of CC .

Let $p \in Sup$, and let $T = cc(I \rightarrow p)$. By Lemma 1, T is consistent, and $T \subset CC$ by construction of CC . We need to show that T is maximally consistent in CC . So let $t' = e' \rightarrow H' \in CC - T$, and let $T' = T \cup \{t'\}$. Since $t' \in CC$, $t' \in cc(I \rightarrow p')$ for some $p' \in Sup$, $p' \neq p$, and $H' = \Pi_{e'}\langle p' \rangle$. We will show that T' is inconsistent.

We first show that on the assumption that T' is consistent, $e' < p$: Assume that there is no history containing both e' and p . Then there must be a splitting point $e'' < e'$, and as in the proof of Lemma 8, we get $e'' \in pcl(I \rightarrow p)$, so $e'' \rightarrow \Pi_{e''}\langle p \rangle \in T$. But as e'' is a splitting point for $H_{\langle e' \rangle}$ and $H_{\langle p \rangle}$, $\Pi_{e''}\langle p \rangle \cap H_{\langle e' \rangle} = \emptyset$, so that T' would be inconsistent. Thus there is a history containing both e' and p , and by Postulate 4, $e' < p$. Now again as in the proof of Lemma 8, $e' \in pcl(I \rightarrow p)$, and thus either $t' \in T$, contradicting our assumption $t' \notin T$, or T' is inconsistent by Lemma 3 (3). \square

(2) ' \supseteq ': Each maximally consistent subset of CC is the set of *causae causantes* of a transition $I \rightarrow p$, $p \in Sup$. We show (a) each consistent subset

of CC determines a supremum of I ; (b) each maximally consistent subset of CC does so uniquely, and (c) a maximally consistent subset of CC is in fact equal to the corresponding set of *causae causantes*.

(a) Let S be a consistent subset of CC , let S_m be the set of transitions maximal in S (by finiteness, $S_m \neq \emptyset$). We show that there is a history in $\bigcap_{t_i \in S_m} H_i$ that contains I :

(a.i) For $t_i \in S_m$, all e_i are pairwise space-like related (by maximality) and belong to one history (Lemma 3 (2)).

(a.ii) Also, for each e_i there is a point $e \in I$ s.t. for all $e' \in I$, if $e' \geq e$, then e' SLR e_i : (α) e_i is a past cause-like locus for some transition $I \mapsto p$, $p \in Sup$. So $e_i < p$, and there is a history h such that e_i and I belong to h and for $h_p \in H_{(p)}$, $h \perp_e h_p$. (β) For all $e \in I$, $e_i \not\leq e$. Assume otherwise. Then $e \in h$ (since $I \subseteq h$). Select $h_p \in H_{(p)}$. Since $I < p$, $e < p$, so $e \in h_p$. But $e_i < e$ now gives $h \equiv_e h_p$, contradicting $h_p \in H_{(p)}$. (γ) There is $e \in I$ such that $e_i \not\leq e$. Otherwise, e_i is an upper bound for I , and $e_i \in h$, so $e_i \geq p' = \sup_h(I)$. But for the $p \in Sup$ from (α), we also had $e_i < p$. By Postulate 4 (using I and $\{e_i\}$), this is impossible. So e SLR e_i . (δ) For $e' \in I$, if $e < e'$, then also e' SLR e_i otherwise, $e < e' \leq e_i$, contradicting (γ).

(a.iii) Since S_m is finite, the intersection of the final sections of I constructed in step (a.ii) is nonempty, call it I' . We have established $\{e_i \mid t_i \in S_m\}$ SLR I' .

(a.iv) At this point we may use Postulate 1 for the primary outcome $\bigcap_{t_i \in S_m} H_i$ of the set $\{e_i \mid t_i \in S_m\}$, and for some primary outcome of I' , to give us a joint outcome and thus a history h containing I' (and thus, I) and contained in all the outcomes H_i of $t_i \in S_m$. Since S_m was the set of maxima, we also have $h \in \bigcap_{t_i \in S} H_i$. So S determines the transition $I \mapsto \sup_h(I)$.

(b) Now let S be maximally consistent, and let $h_1, h_2 \in \bigcap_{t_i \in S} H_i$ with $I \subseteq h_1$ and $I \subseteq h_2$. Set $p_i = \sup_{h_i}(I)$, $i = 1, 2$. We want to show $p_1 = p_2$. Assume $p_1 \neq p_2$. (i) By the prior choice principle, since $p_1 \in h_1 - h_2$, there is a choice point $e < p_1$ such that $h_1 \perp_e h_2$ (i.e., e is maximal in $h_1 \cap h_2$). By transitivity of undividedness, $h_1 \perp_e H_{(p_2)}$. So $e \in pcl(I \mapsto p_2)$. (ii) The transition $t = e \mapsto \Pi_e(p_2)$ is a *causa causans* for p_2 , so $t \in CC$. But $t \notin S$, since t rules out h_1 ($h_1 \notin \Pi_e(p_2)$), and we had $h_1 \in \bigcap_{t_i \in S} H_i$. (iii) Since also $h_2 \in \bigcap_{t_i \in S} H_i$ and $h_2 \in \Pi_e(p_2)$, the set $S \cup \{t\}$ is consistent. This contradicts the assumption that S is maximally consistent in CC .

(c) Let $S \subseteq CC$ be maximally consistent, and let $p \in Sup$ be the uniquely determined supremum of I for which $H_{(p)} \cap \bigcap_{t_i \in S} H_i \neq \emptyset$. We now show that $S = cc(I \mapsto p)$.

(c) ' \subseteq ': Let $t_i = e_i \mapsto H_i \in S$. We have $e_i < p$, and as $e_i \in PCL$, we also have $e_i \in pcl(I \mapsto p)$, and as $H_i = \Pi_{e_i}(p)$, in fact $t_i \in cc(I \mapsto p)$.

(c) ' \supseteq ': Let $t_i = e_i \mapsto H_i \in CC - S$. Then $S \cup \{t_i\}$ must be inconsistent (since S is maximally consistent), and thus it cannot be that there

is $h_p \in H_{(p)} \cap H_i$ (else that history would witness consistency). But then $t_i \notin cc(I \rightsquigarrow p)$. \square

(3) The last part of the proof is simple: since $A_{CC} = \{cc(I \rightsquigarrow p) \mid p \in Sup\}$, the sum in question turns out to be the measure of the unit element $\mathbf{1}_{CC}$, equal to one for any measure μ_{CC} . In detail:

$$\begin{aligned} \sum_{p \in Sup} \mu(I \rightsquigarrow p) &= \sum_{p \in Sup} \mu_{CC}(cc(I \rightsquigarrow p)) = \sum_{T \in A_{CC}} \mu_{CC}(T) \\ &= \mu_{CC}(\cup_{T \in A_{CC}} T) \\ &= \mu_{CC}(\mathbf{1}_{CC}) = 1 \end{aligned} \quad \square$$

Appendix

A Preservation of the ordering of suprema

In the proof of Theorem 2 we had to invoke Postulate 4, which says that the ordering of history-relative suprema is preserved across histories. Here we will show that Postulate 4 is not a consequence of the other Postulates of branching space-times. In the following section we will then show that Postulate 4 is a consequence of branching space-times theory augmented by a notion of ‘same space-time point’.¹³

We will construct a model with two histories, h_1 and h_2 . Let $\langle \mathbb{R}^2, \leq_M \rangle$ be the two-dimensional Minkowskian plane, i.e., the set of points in the two-dimensional real plane with the Minkowskian ordering,

$$(x, y) \leq_M (x', y') \quad \text{iff} \quad (x-x')^2 \leq (y-y')^2 \quad \text{and} \quad y \leq y'$$

This will be our first history, h_1 . In order to have branching, and thus a second history, we add the set

$$F = \{(*, y) \mid y > 0\}$$

which is isomorphic to the positive reals. We now define the ordering \leq on $W = \mathbb{R}^2 \cup F$ to be

$$\begin{aligned} (x, y) \leq (x', y') &\quad \text{iff} \quad (x, y) \leq_M (x', y') \\ (*, y) \leq (*, y') &\quad \text{iff} \quad y \leq y' \\ (x, y) < (*, y') &\quad \text{iff} \quad (x, y) <_M (-y', y') \quad \text{or} \quad (x, y) <_M (y', y') \end{aligned}$$

¹³ The formulation of Postulate 4 is due originally to Weiner ([1997]). He also gave a rather simple model showing that the Postulate is not a consequence of branching space-times alone (Weiner, personal communication). The model presented here was derived independently. It emphasises the fact that histories in a branching space-times model may be radically non-isomorphic.

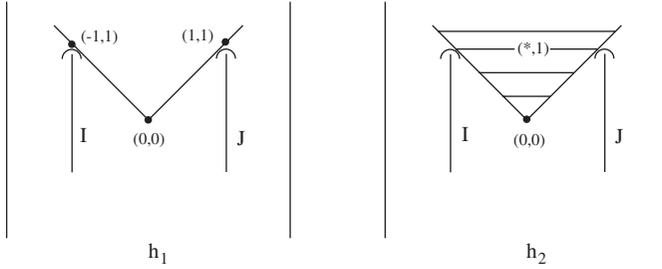


Figure 3. A branching space times model violating Postulate 4. (See text for details.)

Thus, for no (x, y) and $(*, y')$ do we have $(*, y') \leq (x, y)$. In this way, F is pasted into the future light cone above $(0, 0)$. Through the ordering, $(*, y)$ is in some sense identified with the two points $(-y, y)$ and (y, y) that are on the border of the future light cone above $(0, 0)$. The resulting model is $\langle W, \leq \rangle$, pictured in Figure 3. We need to establish that this is a model of branching space-times. Clearly, $W \neq \emptyset$.

Fact: The ordering is antisymmetric.

Proof: Since this is clear for the ordering on \mathbb{R}^2 and on F , we only need to look at the ‘mixed’ case. But if $(x, y) < (*, y')$, it cannot be that $(*, y') < (x, y)$, by the definition of the ordering. □

Fact: The ordering is transitive.

Proof: Select $a, b, c \in W$ such that $a < b, b < c$. If all the three points are in \mathbb{R}^2 or all in F , then $a < c$ follows immediately from transitivity of the respective orderings. So we only need to check two cases: (i) $a = (x, y), b = (x', y'), c = (*, y'')$. Assume that $b < c$ is due to the fact that $(x', y') <_M (y'', y'')$ (for $-y''$, argue analogously). Then $a < c$ follows from the transitivity of the Minkowskian ordering. (ii) $a = (x, y), b = (*, y'), c = (*, y'')$. Assume that $(x, y) < (y', y')$ holds. Then since $y' < y''$, in the Minkowskian ordering, $(y', y') <_M (y'', y'')$, and $a < c$ again follows from the transitivity of the Minkowskian ordering. □

Fact: Let $a = (x, y) >_M (0, 0), b = (*, y')$. Then a and b are incomparable.

Proof: By the definition of the ordering. □

Fact: In $\langle W, \leq \rangle$ there are exactly two histories.

Proof: $h_1 = \mathbb{R}^2$ is a history, since it is upward directed, and by the previous fact, it cannot be extended by elements of F . The set

$$h_2 = (h_1 - \{(x, y) \mid (0, 0) <_M (x, y)\}) \cup F$$

is also upward directed, and cannot be extended by elements of $\{(x, y) \mid (0, 0) <_M (x, y)\}$ either. Again by the previous fact, ‘mixed’ cases are excluded. \square

Fact: The ordering is dense and has infima and history-relative suprema.

Proof: From the properties of the orderings on \mathbb{R}^2 and on F . \square

Fact: The prior choice principle holds. There is exactly one splitting point, $(0, 0)$.

Proof: Let c_1 be a chain in $h_1 - h_2$, i.e., in the future light cone of $(0, 0)$. Then $(0, 0) < c_1$, and $(0, 0)$ is maximal in $h_1 \cap h_2$. Alternatively, let c_2 be a chain in $h_2 - h_1$. Again $(0, 0) < c_2$, since $(0, 0) < (*, y)$ for all $y > 0$. \square

We have thus established that $\langle W, \leq \rangle$ is a model of branching space-times, respecting the Postulates (cf. Belnap [2002]). We now prove that $\langle W, \leq \rangle$ violates Postulate 4.

Fact: In W , there are two initial chains I, J that have different suprema in h_1 , but the same supremum in h_2 .

Proof: Let $I = \{(1, 1 - 1/n) \mid n \in \mathbb{N}\}$, $J = \{(-1, 1 - 1/n) \mid n \in \mathbb{N}\}$, as shown in Figure 3. These chains are in $h_1 \cap h_2$, and they are bounded e.g., by $(1, 1)$ and $(-1, 1)$, respectively, which are their suprema in h_1 . In h_2 , however, $\sup_{h_2}(I) = (*, 1) = \sup_{h_2}(J)$. Along the same lines, one can construct suprema that are space-like related in h_1 , but comparable in h_2 (e.g., set $I' = \{(2, 2 - 1/n) \mid n \in \mathbb{N}\}$). \square

B Branching space-times with space-time points

If we have available a notion of ‘the same space-time point’ in a model of branching space-times, then Postulate 4 holds. We define the notion of ‘same space-time point’ as follows:

Definition 21 (Branching space-times with space-time points) *A triple $\langle W, \leq, \mathbf{S} \rangle$ is a model of branching space-times with space-time points (BST+S) iff $\langle W, \leq \rangle$ is a model of branching space-times and \mathbf{S} is an equivalence relation on W such that*

1. *For each history h in W and for each equivalence class $[s]$, $s \in W$, the intersection $h \cap [s]$ contains exactly one element.*
2. *\mathbf{S} respects the ordering, i.e., for $[s], [s']$ equivalence classes and h_1, h_2 histories, $[s] \cap h_1 = [s'] \cap h_1$ iff $[s] \cap h_2 = [s'] \cap h_2$, and the same for ‘ $<$ ’ and for ‘ $>$ ’.*

In BST+S we can prove that history-relative suprema of initial chains, guaranteed to exist by BST, are at the same location.

Lemma 9 *Let $\langle W, \leq, \mathbf{S} \rangle$ be a model of $BST+S$, let I be an upper-bounded chain in W , and let h_1, h_2 be two histories with $I \subseteq h_1 \cap h_2$. Then $\mathbf{S}(\sup_{h_1}(I), \sup_{h_2}(I))$.*

Proof: Let $s_1 = \sup_{h_1}(I)$, and let $s_2 = [s_1] \cap h_2$ be the space-time point in h_2 that is at the same position as s_1 according to \mathbf{S} . We need to prove that $s_2 = \sup_{h_2}(I)$. First, s_2 is an upper bound for I : take $e \in I$. Since $e \leq s_1$ and \mathbf{S} preserves order and $e \in h_1 \cap h_2$, we have $e < s_2$. It remains to prove that s_2 is the smallest upper bound for I in h_2 . Assume otherwise, then there has to be $s'_2 \in h_2$, $s'_2 < s_2$, and s'_2 an upper bound for I . By the argument just given, $s'_1 = [s'_2] \cap h_1$ is an upper bound for I in h_1 , and from order preservation, $s'_1 < s_1$. This contradicts the assumption that s_1 is the supremum (least upper bound) of I in h_1 . Thus, $s_2 = \sup_{h_2}(I)$: the supremum of I is at the same space-time point across histories. \square

As a corollary, we get that a model of $BST + S$ satisfies Postulate 4.

Corollary 1 *In a model of $BST + S$, Postulate 4 is satisfied*

Proof: Let I, J be two initial chains, and let h_1, h_2 be two histories to which both I and J belong. Then by Lemma 9, the history-relative suprema will be at the same space-time points, and the claim follows by order preservation of \mathbf{S} . \square

With the model of Appendix A we have established the following:

Fact: Not every model of BST can be extended to a model of $BST + S$.

Proof: By the model from Appendix A and by Corollary 1, using contraposition. \square

Acknowledgements

Most of this paper was written while I was a guest of Nuel Belnap's at the University of Pittsburgh in the summer of 2003. I thank him both for his hospitality and for many helpful discussions. Furthermore, I have greatly benefitted from his numerous comments on previous versions of the paper, including the suggestion to 'try the contrapositive' on a previous proof of Lemma 2. I also thank Tomasz Placek for many stimulating discussions and for specific comments. Nuel Belnap, Tomasz Placek and Matthew Weiner have kindly shared their results prior to publication. Support by the Deutsche Forschungsgemeinschaft is also gratefully acknowledged.

*Philosophisches Seminar, LFB III,
Universität Bonn,
Lennéstr. 39, 53113 Bonn,
Germany
Thomas.Mueller@uni-bonn.de*

References

- Aspect, A., Dalibard, J. and Roger, G. [1982]: 'Experimental Test of Bell's Inequalities Using Time Varying Analyzers', *Physical Review Letters*, **49**, pp. 1804-7.
- Bell, J. [1987]: *Speakable and Unspeakable in Quantum Mechanics: Collected Papers on Quantum Philosophy*, Cambridge: Cambridge University Press.
- Belnap, N. [1992]: 'Branching Space Time', *Synthese*, **92**, pp. 385-434.
- Belnap, N. [2002]: 'EPR like "Funny Business" in the Theory of Branching Space Times', in T. Placek and J. Butterfield (eds.), 2002, *Non locality and Modality*, Dordrecht: Kluwer, pp. 293-315.
- Belnap, N. [2003]: 'No Common Cause EPR Like Funny Business in Branching Space Times', *Philosophical Studies*, **114**, pp. 199-221.
- Belnap, N. [2005]: 'A Theory of Causation: Causae Causantes (Originating Causes) as Inus Conditions in Branching Space Times', *British Journal for the Philosophy of Science*, **56**, pp. 221-53.
- Jensen, F. V. [2001]: *Bayesian Networks and Decision Graphs*, Berlin: Springer.
- Pearl, J. [2000]: *Causality*, Cambridge: Cambridge University Press.
- Popper, K. [1982]: *Quantum Theory and the Schism of Physics*, London: Hutchinson.
- Reichenbach, H. [1956]: *The Direction of Time*, Berkeley, CA: University of California Press.
- van Fraassen, B. [1980]: *The Scientific Image*, Oxford: Oxford University Press.
- von Wright, G. H. [1963]: *Norm and Action: A Logical Inquiry*, London: Routledge.
- Weiner, M. [1997]: 'Probability in Branching Space Time', Unpublished, Department of Philosophy, University of Pittsburgh.
- Weiner, M. and Belnap, N. [forthcoming]: 'How Causal Probabilities Might Fit into our Objectively Indeterministic World', *Synthese*.