

How causal probabilities might fit into our objectively indeterministic world*

Matthew Weiner (mweiner@hum.utah.edu)
University of Utah

Nuel Belnap (belnap@pitt.edu)
University of Pittsburgh

Abstract. We suggest a rigorous theory of how objective single-case transition probabilities fit into our world. The theory combines indeterminism and relativity in the “branching space-times” pattern, and relies on the existing theory of *causae causantes* (originating causes). Its fundamental suggestion is that (at least in simple cases) the probabilities of all transitions can be computed from the basic probabilities attributed individually to their originating causes. The theory explains when and how one can reasonably infer from the probabilities of one “chance set-up” to the probabilities of another such set-up that is located far away.

1. Transition probabilities here and there

Imagine two “chance set-ups” that are separated by perhaps millions of miles.¹

When and how could the transition probabilities of two such [1]
chance set-ups be related?

We suggest a rigorous theory of objective single-case event-event transition probabilities that gives a modestly partial answer to question [1]. The theory only makes sense if one takes into account some aspects of the indeterministic and spatio-temporal structure of our world. We shall suggest an answer to [1] under the proviso that there is an absence of Bell-like strange stochastic correlations coming from quantum mechanics. Our chief purpose, however, is not so much to answer [1] as to lay down a general framework for no-nonsense discussions of how causal probabilities might fit into our indeterministic and spatio-temporal world. The basic proposal is that causal probabilities for any transition are inherited exclusively from probabilities ingredient in the

* We thank Tomasz Placek for his significant help via an extensive and instructive correspondence, Thomas Müller for numerous eye-opening conversations and for his corrections to and comments upon a long sequence of drafts, and both Placek and Müller for pre-publication sharing of their results.

¹ Although we do not introduce “chance set-up” as jargon, we always think of one as consisting of an initial event, a family of outcome events, and a probability distribution on those outcomes conditional on the initial.

causae causantes or originating causes of that transition.² We begin with a story involving a simple flip of a coin, so simple that although its telling requires indeterminism, spatio-temporal complications may be downplayed. Later we bring in a second chance set-up that is located far away from the first, at which point we shall need explicitly to consider spatio-temporal relations as well.

1.1. THE CLOCK STORY

The Marshall Fields Clock sits at the corner of State and Randolph in Chicago. Imagine that we are situated there at 3:00 p.m. on a certain Saturday. A trick coin was flipped under the Clock an hour ago, at 2:00 p.m. The altered balance of the coin favored—but did not guarantee—that the coin would land heads-up on the sidewalk. In detail, the chances of the coin showing heads on just that flip were .6 instead of the figure of .5 suggested by the symmetries. As it turned out, however, the coin landed tails, even though the chances of such were only .4. It helps the story if you picture the *Heads*-face of the coin as *Hot pink*, and the *Tails*-face as *Turquoise*.

Perhaps our world is as deterministic as Kant or Hume would have it, so that such talk of “chances” is mere mythology: The coin came up tails, and there’s an end on it. Let us, however, explore the option that our world is in part truly and objectively indeterministic, and in particular let us suppose that the distribution of chances .4 vs. .6 among the Chicago coin-flip outcomes was entirely objective. That is, at any time in the causal past of the 2:00 p.m. flip, there was no settled fact of which outcome would ensue. At those earlier times, there was only the .4 vs. .6 probability distribution on “after 2:00 p.m. the coin will lie heads up and hot pink.” In contrast, after 2:00 p.m. under the Marshall Fields Clock it was a definite matter that the coin lay tails, and that therefore anyone standing under the Clock saw turquoise. There was, that is, a transition under that Clock on that Saturday *from* .4 vs. .6 as to hot pink vs. turquoise *to* determinate or settled turquoise.

² Although (Belnap, 1992) on branching space-times expressed a hope that its particular framework could support a theory of objective probabilities, the foundational ideas bringing that hope to (partial) fulfillment were formulated only five years later, in (Weiner, 1997). This joint essay combines some of the (modal rather than stochastic) ideas of later NB branching-space-times essays (written after yet an additional five years and based in part on (Weiner, 1997)) with the stochastic ideas and results of the aforementioned (Weiner, 1997). It will be obvious that even now our account remains at best decidedly preliminary.

1.2. INDETERMINISM

The Clock story as told so far—and we have not yet added the distant chance set-up—presupposes objective indeterminism in at least the sense that after the flip there are two possible (but incompatible) historical continuations. The situation could be represented in the so-called “branching-time” representation of indeterminism, about which there is much literature. Here we rely on chapters 6, 7, and 8 of (Belnap et al., 2001) (henceforth FF). Crucial to our understanding is the thesis that in the presence of objective indeterminism, we must be careful in our use of the future tense from the perspective of some event. We must take special care to avoid the philosophically clumsy use of the singular term “the future” as if it were a rigid designator. We must distinguish the non-rigid idea of “the future,” which (given objective indeterminism) obviously depends on what occurs next, from the rigid idea of “the future of possibilities,” to use the phrase recommended by FF. This is a way of endorsing the Prior-Thomason suggestion that in the non-rigid use of “the future” in the context of indeterminism, there is a *double* relativization: (1) to a particular momentary event at which the phrase is being evaluated, and (2) to a particular history containing (or a particular historical continuation from) that momentary event. Most philosophers find (1) unproblematic, whereas (2) typically needs explanation; in addition to the chapters cited above, see also, for example, (Belnap, 2002a). The recommended phrase “the future of possibilities” retains relativization (1), but we call it “rigid” since its use no longer involves the more subtle relativization (2).

To speak without tripping ourselves up we need, given objective indeterminism, to consider one event and two (or four) propositions. Our invocation of “event” and “proposition” is intended as firmly based; see §7. Indeed, the whole of the forthcoming analysis of probabilities will use modal-causal ideas found in some earlier essays, as we now indicate.

1-1 CONVENTION. (BST⁹², EPR-fb, NCC-fb, CC, FF)

BST⁹² refers to the modal and causal theory of “branching space-times” developing from (Belnap, 1992) in the following essays:³ EPR-fb

³ We use “BST⁹²” rather than plain “BST” because there are other important workings-out of the general idea of branching space-times each of which one could appropriately call “BST theory” and some of which employ this very acronym. (EPR-fb used “BST-92,” of which BST⁹² is just a reduced form.) Those closest to BST⁹² are (Müller, 2002), and (Placek, 2002). There are also other essays by McCall, Müller, Oksanen, Placek, and Sharlow that explore alternative ways of endowing with probabilities a world of branching space-times. Placek and Müller sometimes use the acronym “SOBST” for “stochastic outcomes in branching space-

refers to (Belnap, 2002b), NCC-fb to (Belnap, 2003b), CC to (Belnap, 2002c), and FF to (Belnap et al., 2001).

The theory relating events to propositions expressing their occurrence is described at length in CC. For a few references to CC, see §7 below.


We label “ e_F ” an appropriate event immediately before hot pink vs. turquoise becomes settled. If you think of O_T as the piece of the world-line of the Clock after the turquoise outcome is settled, and O_H as such a piece after the hot pink outcome is settled, then e_F is a double infimum: $e_F = \inf(O_T)$ and $e_F = \inf(O_H)$. In other words, in BST⁹² theory (and of course it is just a theory), e_F turns out, when idealized, to be a point event, the last point event at which the outcome has not yet been settled.⁴ Then there is the proposition reporting the occurrence of the hot pink outcome and the incompatible one reporting the occurrence of the turquoise outcome. For completeness, we may add the disjunction of these two propositions, which simply says that e_F occurs, and their intersection, which is the inconsistent proposition. These four “outcome propositions” constitute our “probability space,” on which we have laid the probabilities .4 and .6, and of course 1 and 0 in order to satisfy the requirements of abstract probability theory. These probabilities are transition probabilities, conditional on the event e_F occurring. No e_F , no probabilities. They are not “conditional probabilities” that can be calculated in the probability calculus by the standard formula $pr(A/B) = (pr(AB) \div pr(B))$, at least for this reason: No absolute (non-conditional) probability whatsoever is given to the occurrence of e_F .⁵ We only lay the probabilities on the outcomes of e_F , conditional on the event e_F itself. Furthermore, the probabilities concern what occurs *after* the event, e_F , in the causal structure of our world, and for this event you may not (in this theory) substitute some proposition.⁶

times,” which forcefully emphasizes that the target of analysis is in common. The topic is difficult, and needs all the approaches that it can attract. Essays recent to our attention include (Oksanen, 2003), (Sharlow, 1998), (Sharlow, 2003), (Placek, 2003a), (Placek, 2003b), and (Müller, 2003), the last two of which have influenced some aspects of our presentation (as we indicate below). Others can be located by chasing down various references in those just listed as well as in the BST⁹² essays.

⁴ This is a sensible idealization, akin to identifying a billiard ball with its point-like center of mass. It needs to be added, however, that the theory we are going to propose works perfectly well if the initial event of the flip is considered to be a “cloud” of point events. It’s just more complicated.

⁵ This is not to deny that it has one. If it does, however, its probability might well be zero. In any case, such “unconditional” probabilities play no role whatsoever in the story we tell nor in the theory we are developing.

⁶ A modal analog to this theoretical requirement is motivated in CC §5.



The gray regions are two pictures of a *single* set of events.
The white regions picture *two incompatible* sets of events.

Figure 1. Coin-flip under the Marshall Fields Clock

1.3. SPACE-TIME RELATIVITY

We plan eventually to enrich our story with a second chance set-up located on Pluto. Since, however, Pluto is far away from State and Randolph, we cannot do so without imposing the further condition that our world is not only indeterministic, but broadly relativistic in at least the simple-minded sense that the fundamental causal ordering is on local events, even point events, rather than on gigantic world-wide simultaneity slices. Otherwise we would not be able to represent that the flip of the coin is a strictly local matter. Relativity, however, contributes an additional need for care: Given a certain local event, e_1 , there is an objective (frame-independent) *difference* between saying that another event e_2 is (1) in its causal past, or (2) in its causal future, or (3) space-like related to it. We must therefore distinguish the frame-invariant (rigid) idea of “causal future” from the non-rigid idea of “the future,” whose meaning depends on a frame of reference. When we combine indeterminism and relativity, we evidently find that the only *rigid* phrase at our disposal is “the causal future of possibilities,” which depends neither on the frame of reference nor on what occurs next. This is an ugly phrase. We shall nevertheless not try to shorten it (unless from time to time we forget); we need its length to remind ourselves that if we leave out either “causal” or “of possibilities,” our phrase is no longer rigid even though it may sound so. (We must also distinguish the non-rigid (frame-dependent) phrase “the past” from the rigid phrase “the causal past” on grounds of relativity, but there is no additional subtlety added by indeterminism.)

That is about as much as we can do with mere words. To go on, we need to bring in Figure 1. This is a “branching space-times” picture such as occurs in the BST⁹² essays; the conventions governing such

pictures are perhaps best explained in note 23 of EPR-fb. You may use Figure 1 as a help in imagining yourself in different causal situations in our world.⁷

First, suppose you are located under the Clock *before* the flip, say at e_C .⁸ Then, as indicated previously, whether turquoise (tails) or hot pink (heads) is to ensue is not yet a settled matter. There are (to oversimplify) *two* courses of events that are possible for you—with probability distribution as also indicated in the figure: $pr = .4$ for turquoise, $pr = .6$ for hot pink. Complete courses of events that run all the way back and all the way up, as well as all the way out, are called “histories,” and in Figure 1 these two histories are labeled h_1 and h_2 .⁹ The same probability considerations apply if instead of being under the Clock at e_C , you are at some remove, but still in the *causal past* of the flip, say at e_0 .

Second, place yourself still under the Clock, but now *after* the flip, and in particular, (causally) after the less-likely outcome of turquoise has occurred. You can truly say “hot pink was possible before 2:00, and was even the more likely outcome, but it is now a settled matter that turquoise is what occurred.” You might add, “At 1:50 I placed a bet on hot pink at appropriate odds that rendered my bet perfectly fair, but now, shortly after 2:00, it is a settled matter—I can see the turquoise shining with absolute clarity—that I have lost my bet.” Your syntax may become a little tangled up in your effort to be accurate in the context of indeterminism, but it will soon come out all right with the guidance of the Prior-Thomason logic appropriate for those whose world has the indeterministic structure of branching time; we repeat

⁷ Let us note at once that on the one hand, all our pictures will assume that each possible course of events constitutes a Minkowski space-time, but on the other hand, in contrast to some of the other workings-out of BST theory, the postulates and definitions of BST⁹² do not come close to forcing this structure. For instance, it is often said that in Minkowski space-time there is a finite upper limit on the velocity at which an effect can be propagated; but BST⁹² theory is too austere to decide the matter either way.

⁸ Figure 1 registers that the BST⁹² postulates imply that among all the point events at which the outcome of the flip is not yet settled, there is (ideally) a distinguished maximum, namely e_F ; whereas there is no minimum among those at which it is settled that the coin came up turquoise = tails, nor among those at which it is settled that the coin came up hot pink = heads.

⁹ We picture only two histories here for expository simplicity alone. It does not really make sense to attach probabilities to individual histories. Rather, a possible outcome such as turquoise that is fit to receive a probability would be represented as a monumentally large *set* of histories; see §3.


that help can be obtained from chapters 6, 7, and 8 of FF.¹⁰ As before, the same considerations apply if instead of being under the Clock, you are at some remove, but still in the *causal future of possibilities* of the flip, such as at e_3 . Note well that if the coin had come up hot pink, then the event e_3 would not have occurred, since at e_3 it is settled that the coin came up turquoise. BST⁹² theory insists that e_3 in h_1 and e_4 in h_2 are distinct events, and indeed inconsistent, since at one event one could truly say that the coin came up turquoise, whereas at the other one could truly say that the coin came up hot pink. The branching between h_1 and h_2 is located precisely at the point event labeled “ e_F .”

Third, place yourself at an event that is neither in the causal past nor in the causal future of possibilities of the flip; sometimes, thinking of Figure 1, we say that such an event is “in the wings.” In relativity jargon, such events are “space-like related” to the flip. What shall we say about the status of hot pink vs. turquoise at an event that is space-like related to the flip? If we believed that relativity were false and that there is “action at a distance,” then we might be trapped into thinking that the matter is not settled at any such event, such as e_1 , occurring *before* 2:00, while being settled at any event such as e_2 occurring *at or after* 2:00. But this is double talk: Relativity is true, and there is no “absolute” (rigid) sense to saying that an event occurs before or after 2:00. It depends on the so-called frame of reference. BST⁹² theory insists on this: At *either* e_1 *or* e_2 as pictured in Figure 1, since these points are not in the causal future of possibilities of the flip, it is not a settled matter whether the flip will have resulted in hot pink or turquoise. You will have to wait—and 2:00 has nothing to do with anything.¹¹ According to BST⁹² theory, one needs to say that at space-like related events such as e_1 and e_2 (in the wings), the outcome of the flip is not yet settled.¹² From (some interpretations of) special relativity we learn that there is no “action at a distance” (or not much), so that it is wrong to suggest that the effects of the flip coming up turquoise should be transmitted either instantaneously or faster than the fastest signal.

¹⁰ The adaptation of the Prior-Thomason analysis to BST is given in (Müller, 2002), and also in the appendix to CC. There you will also find an account of how “settled” should be used.

¹¹ Note how easily you can see the naive relativity-indeterminism point with the help of the simple BST⁹² diagram of Figure 1. That is not to say, however, that for sophisticated understanding of quantum mechanics one might not need ideas as complicated as “hyperplane dependence” in the sense worked out by (Fleming, 1965), and certainly BST⁹² pictures are nothing like a help in understanding relativistic quantum field theory.

¹² “Not yet” is the correct tense, for the outcome of the flip *will* be settled in the causal future of possibilities of each such space-like related event. See note 10.



The gray regions are two pictures of a *single* set of events.
The white regions picture *two incompatible* sets of events.

Figure 2. Flip under the Clock detected on Pluto

1.4. OVER ON PLUTO

Fourth, use Figure 2 as a help in placing yourself as an investigator traveling on the piece of a world-line on Pluto that we have labeled I_P . It takes light five and one half hours to pass between Earth and Pluto.¹³ Hence, events occurring during a stretch of the life of an investigator on Pluto are all space-like related to the coin flip at State and Randolph. These I_P events are in the wings, neither in the causal past nor in the causal future of possibilities of the flip. For a Pluto event in the *causal past* of the flip (none of these happen to be shown in Figure 2), certainly (given objective indeterminism) one needs to say that the outcome of the flip is not yet determined. And certainly with respect to any Pluto event that lies in the *causal future of possibilities* of the flip, it is a settled matter either that the outcome was turquoise (if the event in question is in history h_1 , for example p_T) or hot pink (if in h_2 , for example p_H). Warning: the diagram shows a horizontal slice for 2:00; but that is intended as relative to the frame of reference in which the Clock and Pluto are at rest. Given just the bare spirit of special relativity, there is no absolute (non-relative) sense to saying that a certain event on Pluto is exactly simultaneous to the two o'clock coin-flip under the Clock. These judgments are not perhaps “intuitive,” but they are inescapable.

The theory of branching space-times runs with these judgments and declares that except when quantum-mechanical EPR-like “modal funny business” threatens (see Assumption 1-2), no matter where you are in the universe, even on far-away Pluto, the outcome of the coin flip

¹³ Do let us imagine that planetary rotations and wanderings have nothing to do with our problem, so that we can pretend that I_P is at rest relative to the Clock situated at State and Randolph.

is not settled as long as you are merely space-like related to it. The outcome becomes settled, according to that theory, only for events in the absolute causal future of possibilities of the flip. Suppose that you have bet on the flip. If you are under the Clock, you need to wait for the flip to take place in order for the bet to be settled; that is, you need to wait until you are in the causal future of possibilities of the flip and, equivalently, the flip is in your causal past. And if instead you are on Pluto, exactly the same thing applies: There, too, you need to wait until you are in the causal future of possibilities of the flip so that you can properly use the causal past tense to report the outcome of the flip. If you keep track of time relatively, according to the frame of reference set up by imagining the Clock at rest, then your bet will be settled at about 7:30, five and half hours “after” the flip. At that point, you will (ideally, of course) see either turquoise or hot pink emanating from the top side of the flipped coin (we are assuming that color shifts are either absent or irrelevant), and the bet can be paid off accordingly.¹⁴

We are finally ready for the story of the promised second chance set-up. In Figure 2 we have marked a track on Pluto as I_P , supposing it to be a piece of the world line of an investigator on Pluto. All of I_P is space-like related to the flip, so that in the course of this track, the outcome of the flip is just as unsettled as it is in relation to events in the causal past of the flip. We have also marked two particular point events in Figure 2 as p_T and p_H . The first marks a possible point event on the investigator’s world-line on Pluto at which she sees turquoise, and the second marks an also-possible event on Pluto at which she sees hot pink.¹⁵ Both are in the causal future of possibilities of the coin flip

¹⁴ This essay is not about frames of reference, but even so it might be helpful to imagine a traveler hurtling by the intersection of State and Randolph on Einstein’s train at a substantial velocity relative to the Clock. Choose the frame of reference that keeps the train at rest (and therefore puts the Clock and Pluto in the same uniform relative velocity relative to the train). Such a traveler, if she remains in h_1 , will *still* see turquoise immediately after the Clock shows 2:00 (Einstein’s coincidence criterion). The change in frame of reference will, in contrast, make a (Lorentz) difference in both distances (e.g. to Pluto) and times (e.g. the temporal interval after 2:00). For this reason, although the absolute *event* of the arrival of turquoise light on Pluto will be quite the same, and although the absolute velocity of light will not vary, nevertheless, since the distance between the Clock and Pluto will be different with reference to the train as at rest, so also the clock time on Pluto read by clocks at rest in the train framework will also be different from 7:30. We mention this, however, only to put it aside, since the deliberations of this essay involve only absolute notions, with (relative) distances and times being used only to help our weak imaginations in understanding illustrations.

¹⁵ “The world line of the investigator” is wrong. Obviously “the investigator” has in our picture a representation that is more like a tree than a line. Is this then some weird “bifurcation” theory? Does the investigator somehow split? No. Bifurcation,

e_F , so that no matter what occurs earlier on Pluto, exactly one of p_T and p_H is going to occur. This is a genuine chance set-up, and one that is far away from State and Randolph. Let us now specialize question [1].

Given I_P , with what probabilities should one expect the [2]
occurrence of p_T vs. the occurrence of p_H ?

Let's put [2] as clearly as possible.¹⁶ In analogy to our treatment of e_F and its outcomes, there is one event and two (or four) outcomes. The event is now the piece of the world line of the investigator, indicated as I_P in Figure 2, lying roughly four billion miles from the Marshall Fields Clock. We are considering the transition from that event, I_P , to its only two possible outcomes (in the story), namely, the proposition that p_T occurs and the proposition that p_H occurs.¹⁷ In strict analogy to e_F , we are asking for transition probabilities, or event-conditional probabilities. (Just as we excluded attaching a probability to e_F , so we are now excluding attaching a probability to I_P .) What is the probability that p_T [p_H] occurs given I_P ? If the investigator on Pluto is in

or splitting of a continuant, means some kind of splitting in a single space-time (amoebas and perhaps personalities do it). To speak carefully in the context of indeterminism without appearing foolish in the course of making an ill-conceived sarcastic remark, one needs to say that the life of the investigator is represented by two world lines, and that those world lines branch *from each other* (it is the set or tree of world lines that branches, not a single world line). In more prosaic terms, for the one investigator on Pluto there are two possible future continuations: She might later see turquoise and she might later see hot pink. In both of these cases, we are speaking of future possibilities *for her*. (Those who think this is a "many investigators" theory are mistaken.) In addition, someone who thinks that it is sensible to suppose that a privileged one of the world lines of the investigator is absolutely "actual" and the others mere "counterparts" can find a critique of this view in chapter 6 of FF.

¹⁶ "Occurs" in [2] is for us a term of art; the truth of the proposition that such and such an event "occurs" is independent of space-time position (so to speak). Technically we represent an occurrence-proposition of an event as a set of histories. For instance, the occurrence-proposition for the point event p_T is the set of histories containing that point event. There is of course also a space-time dependent idea of occurrence, as when we say that p_T will occur, but hasn't occurred yet. For accurate explanation, the dependent idea requires the notion of "settled" due to (Thomason, 1970), adapted to point events. A given proposition (set of histories) is *settled true at* a point event just in case the set of histories representing the occurrence of the point event is contained in (and so "implies") the given proposition. All of this needs careful disentanglement. For a branching-time version, see (Belnap, 2002a) on "double time references." In the meantime, we will try to abide by the convention that the basic idea of "occurrence" is tenseless, and that the tensed use in this essay always implicitly involves the perspectival idea of "settled true at."

¹⁷ As before, there is also their disjunction, which just says that I_P occurs, and their conjunction, which is impossible.

a betting mood and if she wishes her bet to be objectively fair, what odds should she take or offer, while still in the course of traversing I_P , on the prospect that p_T (say) will occur? We hope that it is clear that we have described *two* “chance set-ups,” one at State and Randolph and the other billions of miles away on Pluto. Nevertheless:

Intuitive answer (to question [2]). We can hardly imagine anyone who hears the story and looks at Figure 2 that will not say that conditional on I_P , the chances of p_T occurring are .4, whereas the chances of p_H are .6. [3]

If the chances of turquoise vs. hot pink being sent forth immediately after the flip in Chicago are thus and so, it *must* be that the chances of the Pluto investigator (provided she “finishes” her investigation by traversing all of I_P) receiving turquoise vs. receiving hot pink must be exactly the same. How not?

There is certainly one way that not: Perhaps there is something akin to quantum-mechanical “entanglement” between the two distant chance set-ups. In answering [2] with [3], we will in effect rule out the presence of such weirdness. It is, however, no good proceeding without saying as clearly as we can just what we are ruling out, and we therefore devote a section to this necessary but unrewarding task.

1.5. STOCHASTIC FUNNY BUSINESS

The pure event vocabulary of branching space-times theory, innocent as it is of the language of QM, aspires to capture only two entanglement-like ideas of funny business. Both are paradigmatically distant correlations; the difference is that the correlations of one sort may be described as “modal” since involving only possibility vs. impossibility, whereas the correlations of the second sort involve probabilities and are therefore stochastic instead of modal. The essays EPR-fb and NCC-fb suggested and proved equivalent four different mathematically exact explications of the modal idea, which they called simply “funny business,” but which we here label “*modal* funny business” so as to avoid confusion with the stochastic idea that now assumes its own prominence. One aim of this essay is to explicate in some measure the idea of “*stochastic* funny business.”

It is to be emphasized that at this point the modal idea has already received an explication, whereas explication of the stochastic idea lies ahead. We shall rely on the exact modal idea in explicating the stochastic notion. Here we include that idea by means of an assumption whose exact meaning is given briefly in CC §4.3 and more fully in EPR-fb and NCC-fb.

1-2 ASSUMPTION. (*No modal funny business*) The following are interchangeable formulations of the assumption of no modal funny business: (1) every cause-like locus for an outcome lies in its past (think of superluminal transmissions); (2) immediate outcomes of space-like related initials are always modally independent (think of distant correlations); (3) there is always a prior screener-off (think of Reichenbach's common-cause principle); and (4) there is always a prior common cause-like locus (also reminiscent of Reichenbach).

1-3 DEFINITION. (*BST^{NoMFB}*) We define BST^{NoMFB} as the theory obtained from BST^{92} by adding the no-modal-funny-business assumption 1-2.

As for stochastic funny business, at this point it is only something to be explicated (in §5). Even the target explicandum can only be vaguely and partially indicated. It may help to note at once that modal funny business implies stochastic funny business by identifying impossibility with zero probability. Stochastic funny business, however, can be present even without modal funny business. "Stochastic funny business" is our jargon for what physicists have discovered arises out of peculiar quantum-mechanical "entanglement" of events that are at far remove one from one another. The idea is, however, not quantum-mechanical, as was first made clear by (Bell, 1964). The literature is vast and full of examples and contrary opinions; we only pick out a little piece.

1-4 PARTIAL EXPLICANDUM. (*Stochastic funny business*) Let there be two chance set-ups, simultaneous, far apart, one in Chicago and one on Pluto; see Figure 3. In Chicago there is an initial event, e_1 , of a spin measurement on a particular axis, with possible outcomes \mathbf{O}_{1+} for spin up and \mathbf{O}_{1-} for spin down. On the two immediate outcomes of that measurement, there is a known probability distribution, namely, .4 probability of spin-up (\mathbf{O}_{1+}) and .6 probability of spin-down (\mathbf{O}_{1-}). The situation on Pluto is similar: On Pluto there is an initial event, e_2 , of a spin measurement on a particular axis, with possible outcomes \mathbf{O}_{2+} for spin up and \mathbf{O}_{2-} for spin down. On the two immediate outcomes of that measurement, there is a known probability distribution, namely, .7 probability of spin-up (\mathbf{O}_{2+}) and .3 probability of spin-down (\mathbf{O}_{2-}). You are an investigator on \mathbf{I} , starting somewhere between Chicago and Pluto, and winding up in the causal future of possibilities of both e_1 and e_2 at one of \mathbf{O}_{++} , \mathbf{O}_{+-} , \mathbf{O}_{-+} , or \mathbf{O}_{--} . You try to calculate the probability of a *joint* outcome (one result from each chance set-up) by simple multiplication; for example, your calculation gives $(.4 \times .7) = .28$ as the




Figure 3. Paradigm stochastic funny business

probability that *both* measurements issue in spin-up (\mathbf{O}_{++}). Your experimental results, however, when conjoined with sound judgment, lead you to infer that the joint probability of two spin-up outcomes (\mathbf{O}_{++}) is .3 rather than the calculated .28. That implies that the two chance set-ups are stochastically correlated. Someone suggests that the correlation might be due to a “common cause” influencing each of the two chance set-ups. You respond that this is impossible, since each chance set-up involves a kind of Dedekind-cut-like *immediate* outcome that leaves no room for influences from the past. This makes the envisaged set-ups of this example vastly different from Figure 2, where the key point is that happenings in Chicago “influence” happenings on Pluto. So you have a *distant stochastic correlation without a common-cause explanation*. That’s a paradigm example of *stochastic funny business*. (We shall see later that causal probabilities can “go wrong” in yet other ways, so permit us to emphasize that in our jargon, “stochastic funny business” always connotes distant space-like correlation without a common-cause explanation.)

One would have to survey an enormous literature to be much more helpful, a task we decline in consideration of our lack of expertise. In fact our eventual explication of stochastic funny business will be much sharper than the explicanda, of which readers may make what they will.

1.6. PROBABILITY THEORY IS NOT ENOUGH

Consideration of the threat of stochastic funny business compels (?) us to admit that any answer such as to [2] should be based on a broadly empirical theory. One *may* be able to make it up while sitting in one's philosophical rocking chair, but the analogy to geometry is apt: Such considerations do not remove the theory from the empirical domain. Nevertheless, it is hard to envisage any answer other than [3]. Nor is this essay going to speak differently. We simply emphasize here the following:

You cannot get your transition probabilities for p_T [p_H] given I_P from the probabilities assumed for the transitions from e_F to turquoise [hot pink] by any manipulation of the probability calculus, no matter how sophisticated.

A theory that relates transition probabilities in Chicago to transition probabilities on Pluto cannot be a mere matter of numerical equations. The “chance set-up” in Chicago is one thing, and the “chance set-up” on Pluto is another. Probability theory alone is not going to tell you the relation between the probabilities of two chance set-ups separated by billions or even millions of miles (or even only by a few yards). The answer that the probabilities are the same may be obvious, but how is that answer to be grounded? What, that is, is the theory? Probability theory may be part of it, but it cannot be all of it, because mere probabilities have no way of getting from Earth to Pluto. Nor does it seem plausible that you should need quantum mechanics for such a simple case, nor even a detailed classical physics of what particles do. (Certainly no “standard” interpretation of probabilities rests itself on such detailed physics.) It seems to us that the theory for which we are searching can be a *pure event theory* (no states, no particles, no processes, no language, no minds) that bases its theoretical answer to [2] on nothing more than the indeterministic and relativistic causal orderings that hold among the various events considered. Something needs to be added to probability theory to get us from Chicago to Pluto, but not too much.

That is why it seems good to conjecture that a theory of probabilities built on BST⁹² theory can provide us with a firm foundation. We do not want some metaphorical account such as “probabilities spreading

through space-time.”¹⁸ We are after an exact theory that will get us from the Marshall Fields Clock to the investigator located several billion miles away on Pluto, but, we hope, without any excess baggage. At any rate, that is the presumption on which we make our proposal. When we want a short term for the theory (or at least the topic) of probabilities in BST⁹² (see Convention 1-1) for which we are searching, we use “PrBST.”¹⁹

1.7. METHODOLOGY

Some negative remarks are in order: PrBST, as we shall offer it, is not at all similar to any of the “standard” notions of probability canvassed by philosophers when they survey the history of the interpretation of probability. PrBST is not a “classical” theory (Laplace) since it grounds absolutely nothing on anything like a principle of indifference. It is not a “logical” theory (Carnap) since it has nothing to do with language nor is it intended as *a priori*. PrBST is not a “frequency” theory (Venn) since at bottom it concerns single cases, and for the same reason it is not a “long-run propensity” theory (Popper). Furthermore, it is not a “subjective” theory (Ramsey) since it says nothing about either rationality or anyone’s mind. PrBST is perhaps closest to “single-case propensity” theories (Giere, Fetzer, Mellor; see (Eagle, 2003)), but to the extent that such a theory is supposed to concern propensities of “situations” or of “arrangements of things,” PrBST is but a distant cousin, for such propensity theories are endowed with a far richer vocabulary than that permitted to PrBST. For instance, exactly like Euclidean geometry, PrBST does not come with its own epistemology or relation to norms of rationality. Of course there are e.g. epistemological questions to be raised concerning PrBST, and they are important, but aside from a few scattered informal remarks, epistemology is not a concern of this essay.

Positively PrBST is a theory of objective event-conditional probabilities, as we shall further explain.²⁰ Our stance is that there is

¹⁸ “Probabilities spreading through phase space” would be, for present purposes, worse.

¹⁹ As indicated in note 3, there already exist theories of probabilities in branching space-times, e.g. SOBST. Our aim in constructing PrBST is to add them specifically to BST⁹².

²⁰ We inherit the objective view of probabilities from the (Coffa, 1973) critique of Carnap. Coffa pointed out the overwhelming significance of Carnap’s omission of *truth* as a requirement on inductive explanation. Within this objective perspective, (Salmon, 1984) worked out his notion of causal processes. In BST⁹² theory, processes are replaced by the mathematically more manageable concept of event-event transitions.

no hope of understanding *probabilities* in an indeterminist world of branching space-times without first understanding *possibilities*, since without possibilities there is no space to stand as a support for the probabilities—there is nothing to *be* probable. (We entirely reject the “compatibilist” requirement that any theory must be consistent with determinism.) We will build PrBST on the mathematically rigorous postulates and formally correct definitions of the theory BST⁹² of objective possibilities (see Convention 1-1). There are only two primitives in BST⁹² theory: The causal ordering relation symbolized by “<,” and the set of all “point events,” a set that we call *Our World*. EPR-fb gives all the postulates of BST⁹² in its §2.1, including the “preservation of historical suprema” postulate described (but not named) in its footnote 10. In contrast to the fewness of its primitives, BST⁹² theory involves a great many defined concepts. We do not have the space to repeat here those postulates and definitions, much less repeat the extended motivations by which we make some kind of claim to “material adequacy,” that are to be found in the BST⁹² essays. All we can do is provide, in §7, a mere list of some of the key concepts that are explained in those essays.

2. Idea of event-conditional probabilities

We take as the general idea that we are going to try to elucidate via theory the notion of a probability of a *transition* from an initial event \mathbf{I} to a later outcome event O^* . Thus, using notation explained in §7, we assume that

$$\mathbf{I} \mapsto O^* \tag{4}$$

is a *transition*, which in BST⁹² consists exactly in the assumptions (1) that $\mathbf{I} \mapsto O^*$ is an ordered pair $\langle \mathbf{I}, O^* \rangle$, (2) that \mathbf{I} is an initial event, (3) that O^* is an outcome event, and (4) that O^* is appropriately later than \mathbf{I} , all of which seem required to make objective sense out of the idea of an event-event transition. Each of $e_F \mapsto O_T$ and $I_P \mapsto p_T$ is, for example, an event-event transition. We then add the following syntactic form:

$$pr(\mathbf{I} \mapsto O^*). \tag{5}$$

We may call this an “event-conditional probability.” Our intent is that $pr(\mathbf{I} \mapsto O^*)$ shall be read something like

$$\begin{array}{l} \text{the transition probability of passing from } \mathbf{I} \text{ to the occurrence} \\ \text{of } O^*, \end{array} \tag{6}$$

or

given that you are located at \mathbf{I} , the objective expectation or probability or likelihood of the occurrence of O^* . [7]

As is explicit in [5], the transition with which we endow a probability is from event to event, not from fact to fact. The occurrence propositions for \mathbf{I} and O^* are conceptually important, but that should not cause us to forget that the causal situations of \mathbf{I} and O^* play a role in our understanding. It should nevertheless be noted that our informal readings [6] and [7] treat \mathbf{I} and O^* slightly differently, emphasizing the *location* of \mathbf{I} and the *occurrence* of O^* , although in each case both location and proposition are wanted. The principle reason for inserting “occurrence” on the outcome side of our readings is this: We do *not* want you to confuse an idea such as [7] with a reading such as

given that you are located at \mathbf{I} , the expectation that you yourself will travel to O^* . [8]

Indeed since [8] involves the idea of a continuant, its language falls outside of the purview of PrBST. The transition to the *occurrence* of O^* seems an altogether safer target.

As first examples of our targets, we have the probability (set by our story) of the flip-turquoise transition of Figure 1 and Figure 2,

$$pr(e_F \rightarrow O_T) = .4, \quad [9]$$

and the probability (whose ascription we hope we can justify) of the investigator-turquoise transition of Figure 2,

$$pr(I_P \rightarrow p_T) = .4. \quad [10]$$

Exercise: Read these examples in accord with our suggestions [6] and [7]. With targets such as these in mind, we begin PrBST theory with just a hint of a constraint.

2-1 PRBST POSTULATE. (*Classification of pr*) *pr* is a *natural* partial function. Its domain of definition is confined to event-event transitions $\mathbf{I} \rightarrow O^*$. Its range of values is confined to $\{r: 0 \leq r \leq 1\}$. When $\mathbf{I} \rightarrow O^*$ is in the domain of definition of *pr*, we shall say that *pr*($\mathbf{I} \rightarrow O^*$) is *defined by nature*.

The deductively usable part of the above, such as it is, is the part omitting “natural” and “by nature.” We include the adjective and the phrase in spite of their lack of deductive content so as informally to indicate our intended application. For example, it is our intention

that [9] and [10], different though they are, shall each be construed as reporting a fact of nature.

3. Basic and basic^β probabilities

Before we move over to Pluto, we should articulate what is going on under the Clock. In BST⁹² we laid great reliance on the notion of a “basic transition” $e \mapsto O$ from a point event e to one of its basic chain outcomes O , or to what is equivalent, to $\Omega_e \langle O \rangle \in \Omega_e$, or indeed to an appropriate basic propositional outcome $H \in \Pi_e$; see §7. *Basic probabilities* are just probabilities of basic transitions, in whichever guise. [9] reports a basic probability since $e_F \mapsto O_T$ is a basic transition. There are also what are in effect boolean combinations of certain basic transitions; since one needs a boolean algebra underlying a probability distribution, this is hardly surprising.

3-1 DEFINITION. (*Algebras Ω_e^β and Π_e^β of basic^β outcomes*) Ω_e^β may be defined as the complete atomic boolean algebra of basic^β outcome events that results from Ω_e by taking all subsets of Ω_e . Ω_e is its 1, and \emptyset is its 0.²¹ A member of Ω_e^β is a set of pairwise incompatible scattered outcome events, hence each $O \in \Omega_e^\beta$ fits the definition of a disjunctive outcome event (see §7 for both “scattered” and “disjunctive” outcome events).

The propositional analog is Π_e^β , which is defined as the complete atomic boolean algebra that results from Π_e by adding unions of all subsets of Π_e . Evidently each member of Π_e^β is a set of histories to which e belongs (= in which e occurs), the 0 of this algebra is \emptyset , and $\{h: e \in h\}$ is its 1.

We will systematically use the superscript “^β” in order to indicate boolean-related concepts.²² Thus, a *basic^β outcome* of e is a member of Ω_e^β or Π_e^β , and a *basic^β transition* has the form $e \mapsto O$ with $O \in \Omega_e^\beta$, or $e \mapsto H$, with $H \in \Pi_e^\beta$.

The step from basic to basic^β transitions is small; it is a much more substantial jump to see that [10] reports a probability that is natural but neither basic nor basic^β, since $I_P \mapsto p_T$ is a transition, but neither basic nor basic^β. We delay that step.

²¹ We have let in \emptyset as a basic^β outcome event, not because we like it, but so we do not constantly have to remember to leave it out.

²² We ask you to put up with this ugly usage for a while in order to mark the sharp analytical divide between basic^β and non-basic^β probabilities that is crucial to this essay.

Nature is boss, and how much she constrains basic probabilities is how much she constrains them, as illustrated in the following.

3-2 EXAMPLE. (*Three basic outcomes*) Suppose $\Omega_e = \{\mathbf{O}_1, \mathbf{O}_2, \mathbf{O}_3\}$, so that the \mathbf{O}_i are pairwise disjoint basic (scattered) outcomes of e exactly one of which must occur if e occurs. We are considering the three basic transitions t_1, t_2 , and t_3 , where $t_i = (e \rightsquigarrow \mathbf{O}_i)$. Here is where boolean algebra appears: We cannot avoid considering as well the set of all eight basic ^{β} transitions $e \rightsquigarrow \mathbf{O}$, where \mathbf{O} is any member of the boolean algebra Ω_e^β , interpreted as a transition to a disjunctive outcome of e .

Philosophy must allow that nature could tell us any of the following. (1) There is no natural sense to be made of comparing the likelihoods of these transitions. (2) t_1 is more likely than t_2 , but there are no natural numerical comparisons. (3) t_1 is twice as likely as t_2 , but there is nothing to say about t_3 . (4) There is a natural probability distribution on these transitions, by which we mean that nature permits us to interpret her by attaching numbers n_i ($1 \leq n \leq 3, 0 \leq n_i \leq 1$) respectively to each of t_1, t_2 , and t_3 in such a way that they are the basis of a Kolmogorov probability distribution on the set of transitions to the boolean algebra Ω_e^β of all basic ^{β} outcomes of e ; for instance, it must turn out that since \mathbf{O}_1 and \mathbf{O}_2 are inconsistent, $pr(e \rightsquigarrow \{\mathbf{O}_1, \mathbf{O}_2\}) = n_1 + n_2$. (The sequence (1)–(4) reminds us of the idea of (van Fraassen, 1980) (p. 198) that “probability is a modality, it is a kind of graded possibility.”)

Someone is going to ask about the “interpretation” of these numbers in case (4). They are transition probabilities, or more exactly, event-conditional transitional probabilities. No number is given to the occurrence of e , nor is any number assigned absolutely to the occurrence of \mathbf{O}_i . Instead, nature has fixed things so that it is settled at e that exactly one member of Ω_e will occur in the immediate future of e , and it is settled at e that the probability distribution on these outcomes is truly given by the numbers n_i , and that the distribution on the engendered basic ^{β} outcomes comes by simple addition. It seems to us that although we may not limit nature *a priori*, there is no choice about this in the following sense: Otherwise we are not speaking of probabilities.

If you wish to use these probabilities to guide you as to how you ought to behave (assuming you are aware of nature’s probability distribution), then you should use them as conditional advice: If you consider what can occur immediately after e (perhaps because you are about to reach e), and if you wish your expectations to conform to nature, you

should expect each basic ^{β} outcome to the degree indicated by the sum of the n_i attached to its members, given e .²³

Someone else is going to ask about the epistemology of these basic transition probabilities. In one sense this is a genuine problem, and a problem to which we have no contribution to make. As we stated earlier, our strategy is to develop the theory without being distracted by subtle (but reasonable) epistemological questions. This is a matter of postponement rather than neglect, invoking the assumption that sometimes trying to be “too epistemological” interferes with the development of useful theory. Still, we are as sure as we can be that the epistemology of any broadly *a posteriori* theory such as ours needs to have a broadly empirical element, and doubtless one involving appeal to something like repeated experiments or observations carried out by persons of good scientific judgment, and of a mixture of deductive and non-codifiable inferences from this empirical base, presumably guided by theory, carried out by the same or by other persons of equally good judgment. This is modest common-sense empiricism that falls far short of an “epistemology.”

For yet someone else, the question may be founded on a belief that the epistemology of determinism is somehow easier than the epistemology of indeterminism, or that the epistemology of the single case is more difficult than the epistemology of the general case. Philosophers do differ radically and sometimes unpersuadably in such beliefs, and we have ours, but the matter seems to us irrelevant. We have not tried to support our theory in terms related to such considerations.

This theory can plausibly serve as a solid foundation (we do not say “the” solid foundation) for attributions of probability, if any event-event attributions are to be had. The idea comes in steps. The first step is the PrBST account of “basic probabilities.”

3-3 PRBST POSTULATE. (*Basic probability*) pr is defined by nature for at least one basic transition. That is, there is a point event e and a scattered outcome event $\mathbf{O} \in \Omega_e$ such that $pr(e \rightarrow \mathbf{O})$ is defined by nature.

[9] is intended as an illustration of PrBST postulate 3-3. This postulate says hardly anything, but setting it down helps to distinguish its force from that of the following, which is a mere definition.

3-4 PRBST DEFINITION. (*pr defined for e*) pr is defined for $e \leftrightarrow_{df}$ (1) $pr(e \rightarrow \mathbf{O})$ is defined by nature for each $\mathbf{O} \subseteq \Omega_e$, and (2) Ω_e is countable.

²³ This advice is like that counselling a bowler to knock down all ten pins. In the words of (Weiner, 2003), such advice is “effective” but not necessarily “doable.”

(We will not distinguish the basic transition $pr(e \rightarrow \mathbf{O})$ from the basic ^{β} transition $pr(e \rightarrow \{\mathbf{O}\})$).

Why do we make “defined for e ” in PrBST definition 3-4 a definition? The perhaps falsely subtle point is that we do not assume or assert that if nature defines $pr(e \rightarrow \mathbf{O})$ for some $\mathbf{O} \in \Omega_e$ then it inevitably defines $pr(e \rightarrow \mathbf{O})$ for all. PrBST doesn’t claim that much about nature. In the same spirit, we refrain from assuming that Ω_e is invariably countable. Still, as far as we can now see, we shall be able to treat pr as a probability only when pr is defined for e in the sense of PrBST definition 3-4. The reason is that so much is needed to make sense of requiring pr to satisfy the following (standard) principle indicating how basic ^{β} probabilities arise from basic probabilities. (Well, the countability requirement could be weakened at the expense of some mathematics that we judge too heavy to be worthwhile in this introductory investigation.)

3-5 PRBST POSTULATE. (*Basic ^{β} Kolmogorov*) If pr is defined for e , then $pr(e \rightarrow \mathbf{O})$ for $\mathbf{O} \subseteq \Omega_e$ satisfies the standard Kolmogorov properties: (1) $pr(e \rightarrow \mathbf{O}) \geq 0$, (2) $pr(e \rightarrow \Omega_e) = 1$, and (3) $pr(e \rightarrow \mathbf{O}) = \sum_{\mathbf{O} \in \mathbf{O}} pr(e \rightarrow \mathbf{O})$.

We conceive PrBST postulate 3-5 as a broadly empirical assumption, although as we indicated, our imaginations are not rich enough to bring into focus any alternatives while still considering pr anything like a probability. Consider Example 3-2. We can (more or less) imagine any of the alternatives (1)-(3), but we cannot (we report) imagine that nature puts a distribution on all three outcomes, as in alternative (4), but whose three numbers signify, for instance, that *each* of \mathbf{O}_1 , \mathbf{O}_2 , and \mathbf{O}_3 is more likely than not. That is the sort of thing we mean by implying via PrBST postulate 3-5 that the three numbers “must” sum to 1, even though of course numbers can be used as codes of likelihoods in heroically many ways. In any event, since imaginings and the lack thereof do not export, we do not intend these remarks as argumentative. The upshot is that according to PrBST, basic ^{β} probabilities are straightforwardly determined from basic probabilities in the simplest possible fashion.

In closing this section on basic and basic ^{β} transition probabilities, it is perhaps not useless to remark that PrBST postulate 3-5 says nothing at all about whence these probabilities come, nor about connections between basic probabilities $pr(e_1 \rightarrow \mathbf{O}_1)$ and $pr(e_2 \rightarrow \mathbf{O}_2)$ for distinct e_1 and e_2 .

4. Causal theory of non-basic^β probabilities

Example target [10] $pr(I_P \rightsquigarrow p_T) = .4$ is *not* a probability of a basic^β transition, an observation that drives most of the remainder of what we have to say. Specifically, a transition $\mathbf{I} \rightsquigarrow O^*$ that is non-basic^β is so in virtue of one of two qualities: Either \mathbf{I} is not a single point event e (i.e., not the unit set of such) or we do have a case of $e \rightsquigarrow O^*$, but O^* is not a basic^β outcome of e as a member of Ω_e^β .²⁴ In either case, the postulates of §3 offer no help, and thus give us no information making sense out of [10]. We first postulate that help is somehow to be found.

4-1 PRBST POSTULATE. (*Non-basic^β probability*) We assume that sometimes there is a natural fact of the matter as to a numerical probability for a non-basic^β transition $pr(\mathbf{I} \rightsquigarrow O^*)$.

For example, in our illustration the non-basic^β probability, [10] $pr(I_P \rightsquigarrow p_T) = .4$, is a fact of nature. We shall say that such a probability (whether basic, basic^β or non-basic^β) is *real*. In speaking of “real” probabilities, we mean to be calling to mind the notion of “possibilities based in reality” of (Xu, 1997). Without, as usual, saying anything epistemological or linguistic, our intent is that when provided, $pr(\mathbf{I} \rightsquigarrow O^*) = r$ is a perfectly objective fact. We certainly think that about [10].

We are going to propose that in well-behaved cases there is a simple relationship between basic probabilities and non-basic^β probabilities that relies exclusively on the causal order. Such is a central idea of this essay. It is part of the proposal that not all cases are “well-behaved,” with stochastic funny business (1-4) being a known way that things can go wrong. This theory of distant probabilities amid well-behaved phenomena is both interesting (because it goes beyond mere numerical calculations with probabilities) and also rigorous. The theory proposes that not just any basic probabilities will do. Picture yourself as the investigator on Pluto so that the non-basic^β transition in which you are interested is $I_P \rightsquigarrow p_T$. “Here-now” is I_P and “here-later” is p_T . The kernel idea might be put (very) roughly and (somewhat) inaccurately as follows: You should consider probabilities of basic transition events whose initials are “over there and back then” and whose outcomes are consistent with the target non-basic^β outcome here-later. That recipe, when spelled out with precision, will send you “over there and back then” to the basic probability $pr(e_F \rightsquigarrow O_T)$ of the transition from flip to turquoise = tails. Precision, as always, costs complexity. The suggested theory linking the basic transition $e_F \rightsquigarrow O_T$ to the non-basic^β transition

²⁴ Again, we can always count members of Π_e^β as basic^β since conceptually interchangeable with members of Ω_e^β .

$I_P \mapsto p_T$ builds on the BST⁹² theory of *causae causantes*, which we now review.

4.1. *Causae causantes*

There is a modal-causal story about causation for such event-event transitions, namely, the theory of *causae causantes* or originating causes (we use these as synonyms) of CC. We pursue the thought that the relation between *causae causantes* considered as causal and the caused transitions $\mathbf{I} \mapsto O^*$ for which they serve as causal influences can provide a foundation for a happy suggestion as to how probabilities fit into our world both under the Clock and on Pluto.

It needs emphasis that the theory of *causae causantes* presented in CC is only known to work given the assumption that there is no modal funny business (1-2). This limitation on our suggestion concerning probabilities is so heavy that in order to keep it in mind, we invented a name for the result of adding to BST⁹² theory the constraining assumption of no modal funny business, namely BST^{NoMFB} (Definition 1-3). In that axiomatic theory we reach the idea of a *causa causans* in three steps. Then we can finally start talking about non-basic^β probabilities.

The theory of *causae causantes* of transitions $\mathbf{I} \mapsto O^*$ bifurcates according to whether O^* is a scattered outcome event \mathbf{O} (all of the pieces of which can begin in a single history) or a disjunctive outcome event \mathbf{O} (reified as a set of pairwise incompatible scattered outcome events). The second part, however, can be left implicit, so that we shall need to be explicit only about *causae causantes* of $\mathbf{I} \mapsto \mathbf{O}$.

4-2 BST⁹² DEFINITION. (*Past causal loci and causae causantes of $\mathbf{I} \mapsto \mathbf{O}$*)

1. In BST^{NoMFB} one can define the idea of a *past causal locus* of an event-event transition $\mathbf{I} \mapsto \mathbf{O}$, to wit, e is a past causal locus of $\mathbf{I} \mapsto \mathbf{O}$ just in case e , being appropriately in the causal past of the outcome event \mathbf{O} , is crucial to the transition in the following sense: There is some history h in which the initial \mathbf{I} finishes such that e is a last point at which both h and \mathbf{O} are possible. So \mathbf{I} finishes in h , and h and \mathbf{O} are compossible up to and including e , but immediately after e , no matter what happens, at least one of h and \mathbf{O} becomes impossible. We write “ $pcl(\mathbf{I} \mapsto \mathbf{O})$ ” for the set of past causal loci of $\mathbf{I} \mapsto \mathbf{O}$.
2. We can put this in a different way. One can prove in BST^{NoMFB} the crucial fact that when e is a past causal locus of $\mathbf{I} \mapsto \mathbf{O}$, exactly one of its several basic outcomes is consistent with \mathbf{O} . We call this

one “ $\Omega_e\langle\mathbf{O}\rangle$,” the occurrence proposition of which might be read as “ \mathbf{O} remains possible at least in the immediate future of e .” So when $e \in \text{pcl}(\mathbf{I} \rightarrow \mathbf{O})$, what occurs immediately after e is either that \mathbf{O} is rendered impossible, or, via its uniquely determined basic outcome $\Omega_e\langle\mathbf{O}\rangle$, that \mathbf{O} is kept possible at least for a little while.

3. We are thereby led to a good definition of a *causa causans* of $\mathbf{I} \rightarrow \mathbf{O}$ in $\text{BST}^{\text{NoMFB}}$: When e is a past causal locus of that transition, then the uniquely given basic transition $e \rightarrow \Omega_e\langle\mathbf{O}\rangle$ is defined as a *causa causans* of the transition $\mathbf{I} \rightarrow \mathbf{O}$.

4-3 EXAMPLE. (*Causa causans from the Clock to Pluto*) The under-the-clock transition $e_F \rightarrow O_T$, or, equivalently, $e_F \rightarrow \Omega_{e_F}\langle O_T \rangle$, as extracted from Figure 2, is an example of a basic transition. Also the Plutonic transition $I_P \rightarrow p_T$ is an example of a non-basic ^{β} transition. It is furthermore easy to see that the basic transition $e_F \rightarrow \Omega_{e_F}\langle O_T \rangle$ is a *causa causans* of the non-basic ^{β} transition $I_P \rightarrow p_T$. The reason is this: Since e_F is in the causal past of p_T , and since what occurs immediately after e_F can render p_T either possible or impossible, it has to be that $e_F \in \text{pcl}(I_P \rightarrow p_T)$. Therefore $e \rightarrow \Omega_e\langle p_T \rangle$ has to be a *causa causans* of $I_P \rightarrow p_T$. The final piece is this: $e_F \rightarrow \Omega_{e_F}\langle O_T \rangle$ is precisely identical to $e_F \rightarrow \Omega_{e_F}\langle p_T \rangle$. That is, the basic transition from e_F that guarantees the beginning of O_T is exactly the same basic transition that keeps p_T possible (at least for a while). That is exactly why $e_F \rightarrow O_T$ is a *causa causans* of $I_P \rightarrow p_T$.

Penultimately, we intend by what *isn't* shown in the simple Figure 2 that the flip-to-turquoise transition under the Clock, namely $e_F \rightarrow O_T$, shall be the *only causa causans* of the investigator-to-turquoise transition $I_P \rightarrow p_T$ on Pluto. (Although the particular example features but a single *causa causans*, $\text{BST}^{\text{NoMFB}}$ theory realistically requires only that it makes sense to speak of the set of all of them, however many there may be.) And finally we ask your agreement that all these statements are fit to be made in the pure causal-ordering language of $\text{BST}^{\text{NoMFB}}$, and are provable from the axioms thereof, so that (except for communication rather than proof) we shall not be understood as relying on untrustworthy pictures.

4.2. NON-BASIC ^{β} PROBABILITIES VIA *causae causantes*

Our causal-stochastic suggestion is that the probability of an arbitrary transition $\mathbf{I} \rightarrow O^*$ depends on the basic transitions that stand as its causes in the sense of its *causae causantes*. That's the causal part of the proposal, and is certainly its true essence. The simplest possible

numerical suggestion is this: If O^* is a non-disjunctive (scattered) outcome event \mathbf{O} just multiply, and if O^* is a disjunctive outcome event \mathbf{O} , just add. This recipe is too easy to be universally true; already stochastic funny business forms a counterexample. Let us nevertheless lay out the simplest suggestion in detail so that we may consider something definite. We shall divide the simplest suggestion into “postulates” and “principles.” The division is based on attitude: The *postulates* we expect to carry over to not-so-simple cases, whereas the *principles* may drop by the wayside.

4-4 POSTULATE. (*Nothing but causae causantes*)

1. Assume that $\mathbf{I} \rightarrow \mathbf{O}$ is a non-basic ^{β} transition from an initial event to a scattered (non-disjunctive) outcome event. $pr(\mathbf{I} \rightarrow \mathbf{O})$ is defined (by nature) iff pr is defined (by nature) for every member e of $pcl(\mathbf{I} \rightarrow \mathbf{O})$ (in other words, for each *causa causans* of $\mathbf{I} \rightarrow \mathbf{O}$).
2. Assume $\mathbf{I} \rightarrow \mathbf{O}$ is a transition from an initial event to a disjunctive outcome event. $pr(\mathbf{I} \rightarrow \mathbf{O})$ is defined (by nature) iff $pr(\mathbf{I} \rightarrow \mathbf{O})$ is defined (by nature) for every $\mathbf{O} \in \mathbf{O}$.
3. When $pr(\mathbf{I} \rightarrow O^*)$ is defined, nothing counts except nature-given stochastic features of its *causae causantes*—including the possibility that one may need to take into account not only probabilities of individual *causae causantes*, but also probabilities of certain sets of them, taken as operating jointly. In other words, the probability of a non-basic ^{β} transition from here-now to here-later depends entirely on stochastic features of basic transitions located “over there and back then.”

We devote a paragraph to a hesitant discussion of the epistemological status of Postulate 4-4(3). We may isolate the topic by considering our particular coin-flipping example, which specifies but a single *causa causans*, and for this reason involves no hint of either multiplication or addition. It may then seem conceptually difficult to make sense of the idea that the probability at State and Randolph of the coin’s coming up heads should be different from the probability on Pluto of the coin’s being seen to come up heads. Nevertheless, it would appear that the evidential basis for the Pluto-based judgment that $pr(\mathbf{I} \rightarrow \mathbf{O}) = r$ may in principle be different from the evidential basis for the Chicago *causa causans* probabilities $e \rightarrow \Pi_e(\mathbf{O}) = r'$ that go into the calculation of [11]. For example, it could turn out that although repeated experiments at State and Randolph confirmed (for persons of good judgment) the turquoise vs. hot pink odds as .4 vs. .6, the repeated experience of investigators on Pluto told a quite different story.

Since involving but a single *causa causans*, this would concern neither multiplication nor addition. What would be at issue is the very idea of a connection between probabilities and the causal order. In such a case persons of good judgment might legitimately decide that the tight connection between probabilities of *causae causantes* and probabilities of non-basic ^{β} transitions asserted by Postulate 4-4(3) must be given up in a way roughly analogous to the usual verdict in the case of Bell-like phenomena.²⁵ That possibility would seem enough to give Postulate 4-4(3) some measure of empirical content.

4.3. JUST MULTIPLY AND ADD

We next put some flesh on the rather bony Postulate 4-4 by considering how one might actually calculate the probability of a non-basic ^{β} transition to a scattered (non-disjunctive) outcome event, thus combining a little arithmetic with the causal order of *Our World*.

4-5 PRINCIPLE. (*Just multiply*) Assume that $\mathbf{I} \rightarrow \mathbf{O}$ is a non-basic ^{β} transition from an initial event \mathbf{I} to a non-disjunctive outcome event \mathbf{O} , and that $pr(\mathbf{I} \rightarrow \mathbf{O})$ is defined (by nature). Assume further that $pcl(\mathbf{I} \rightarrow \mathbf{O})$ is finite. Then *just multiply*:

$$pr(\mathbf{I} \rightarrow \mathbf{O}) = \prod_{e \in pcl(\mathbf{I} \rightarrow \mathbf{O})} pr(e \rightarrow \Omega_e(\mathbf{O})). \quad [11]$$

Observe that we postulate the identity [11] only when the set of *causae causantes* of $\mathbf{I} \rightarrow \mathbf{O}$ is finite. We impose the finiteness limitation because we are unskilled in the use of infinite multiplications, and must therefore leave it to others to remove the finiteness requirement by means of a better theory. The *core* idea of the principle is that we can link the probability of a non-basic ^{β} (“caused”) transition $\mathbf{I} \rightarrow \mathbf{O}$ to the basic probabilities associated with its *causae causantes* by taking into account the contribution of each and every one.

For non-disjunctive outcomes \mathbf{O} , the chief theory of CC makes plausible the “just multiply” principle. There it is proved in BST^{NoMFB} theory that in the case of a non-disjunctive outcome \mathbf{O} , if one takes the set of all *causae causantes* $e \rightarrow \Pi_e(\mathbf{O})$ of $\mathbf{I} \rightarrow \mathbf{O}$, one finds that they form a set of “inns” conditions: Typically each is an *insufficient* condition, and always each is a *necessary* condition, each is *non-redundant*, and jointly they are *sufficient*. The “non-redundancy” clause is critical. It means that each *causa causans* of $\mathbf{I} \rightarrow \mathbf{O}$ has its own separate

²⁵ Because of our choice of jargon, that would not count as “stochastic funny business.” It would, however, fall under the much looser idea of ill-behaved stochastic phenomena, i.e., phenomena that would give scientists a good deal to puzzle over. That, in our view, is the extent of the rough analogy with the Bell situation.

contribution to make, a contribution that cannot be taken up by the others. If you omit even one *causa causans* of $\mathbf{I} \mapsto \mathbf{O}$, the outcome of the “effect” transition wouldn’t occur even if its initial were to occur. *That*, we think, is what makes the rule of simple multiplication plausible, for it would certainly not make sense to multiply by the probability of some condition that, even though a necessary condition in the usual sense, is redundant and could be omitted, its work being taken up by other conditions. Nevertheless, one doesn’t have to read very far into the Bell literature to come to believe the principle false. We delay a closer analysis of its failure.

Next, suppose we wish to link the probability of a transition to a disjunctive outcome event \mathbf{O} to certain basic probabilities. CC defines disjunctive outcome events as sets of pairwise-incompatible scattered outcome events, the thought being that such an event occurs just in case exactly one of its members occurs. We suppose there is more room for applying the idea of a disjunctive outcome event than that of a non-disjunctive one; although on the one hand it is very likely easier to think about the transitions that causally influence a particular occurrence localized in *Our World*, on the other it is generally more interesting to loosen up one’s consideration by thinking of the many ways in which “the” occurrence might have happened. In any event, it does seem straightforward to figure out what the probability for a transition with a disjunctive outcome event \mathbf{O} “should” be. Keep in mind that \mathbf{O} is a set of pairwise incompatible scattered outcome events, and it is proven in CC that its *causae causantes* form a set of *inus* conditions in Mackie’s sense: (typically) *insufficient* but *non-redundant* part of a (typically) *unnecessary* but *sufficient* condition. The Mackie recipe does not include that the various sufficient conditions should be pairwise incompatible, something that seems required in characterizing probabilities, as follows.

4-6 POSTULATE. (*Just add*) Assume that $\mathbf{I} \mapsto \mathbf{O}$ is a transition from an initial event \mathbf{I} to a disjunctive event \mathbf{O} , and that $pr(\mathbf{I} \mapsto \mathbf{O})$ is defined (by nature). Assume further that $pcl(\mathbf{I} \mapsto \mathbf{O})$ is countable. Then, recalling that members of \mathbf{O} are pairwise incompatible, *just add*:

$$pr(\mathbf{I} \mapsto \mathbf{O}) = \sum_{\mathbf{O} \in \mathbf{O}} pr(\mathbf{I} \mapsto \mathbf{O}). \quad [12]$$

This postulate shares two features with PrBST postulate 3-5. First, its range of application is limited by countability, a limitation that could presumably be weakened at the cost of introducing some additional mathematical complications that, we think, do not really touch the heart of our suggestion. The requirement of countability is intended as a limit on the applicability of PrBST theory rather than anything

like an empirical assumption. Second, this postulate seems required in order to say that we are speaking of probabilities rather than of some more general idea of “graded possibilities.” Perhaps not, but that is how it seems to us, which is why we label it a postulate. In any case, this second feature does not convert the postulate from empirical to non-empirical, since we do not need or wish to rule out the possibility that “addition” is inappropriate to the phenomena.

5. Four principles of causal probabilities

Following the spirit of the suggestions of (Placek, 2003b) and the detailed developments of (Müller, 2003), it is possible to analyze rather more closely the possible failures of the just-multiply and just-add principles. We can see that there are four distinct types of failure that threaten. Adequate clarity on the four ideas, when expressed in the fashion of Müller’s workings-out, needs additional careful conceptual elaboration that is inappropriate to this essay. In particular, one wants at least the idea of imposing a causal order on individual *causae causantes* and so derivatively on sets of them, the idea of consistency of sets of *causae causantes*, and the idea of taking an entire consistent set T of *causae causantes* as a proper argument for pr , giving the probability of the “joint transition” T . In these terms we can state the fundamental *causae causantes* postulate, and then we can see that three familiar stochastic properties come into their own *in a causal form*, which we here state without full precision.

1. *Causae causantes property.* Consider any transition $\mathbf{I} \rightarrow \mathbf{O}$ from an initial event \mathbf{I} to a scattered outcome event \mathbf{O} . Let T be the set of all *causae causantes* of $\mathbf{I} \rightarrow \mathbf{O}$. Then $pr(\mathbf{I} \rightarrow \mathbf{O}) = pr(T)$.

This is intended as a more precise version of Postulate 4-4(3). When and if it fails to apply, the ideas of this essay are of no known use.

2. *Marginal property (Müller formulation).* Take a consistent set T of basic transitions and a point event e . Assume that for each $\mathbf{O} \in \Omega_e$, the basic transition $e \rightarrow \mathbf{O}$ is causally maximal in the set $T \cup \{e \rightarrow \mathbf{O}\}$. Then $pr(T) = \sum_{\mathbf{O} \in \Omega_e} pr(T \cup \{e \rightarrow \mathbf{O}\})$.

This property, although stated in causal terms, is not “very” causal because of the provable (Müller) connections between causal maximality and consistency. It is plausibly required just to make sense out of calling pr a probability, just as we suggested for the other “summation” principles PrBST postulate 3-5 and Postulate 4-6 that were offered above.

3. *Markov property (Müller formulation)*. Let $T_1 \cup T_2$ be consistent and let every basic transition in T_1 causally precede every basic transition in T_2 . Then the probability of T_1 is independent of that of T_2 : $pr(T_1 \cup T_2) = (pr(T_1) \times pr(T_2))$.

We might call the Markov property “vertical multiplication.” This is a highly plausible candidate for universality, but it appears more “empirical” than the marginal property in the sense that it seems to go beyond the mere requirement of calling pr a probability. Perhaps it is required by the very nature of the combination of causal ordering and probability; and perhaps not.²⁶ We are reluctant to pretend that we presently have clear insight into the matter.

4. *The no-stochastic-funny-business property (or the factoring property; after both Placek and Müller)*. Let T be consistent and let each basic transition in T be space-like related to every other basic transition in T . Then the probability of each member of T is independent of that of the others: $pr(T) = \prod_{t \in T} pr(t)$.

We might call the no-stochastic-funny-business property “horizontal multiplication.” Its failure is our candidate explication of “stochastic funny business” in the sense of Partial explicandum 1-4. If that is correct, then we know (or think we know) that the principle is not universal: There is stochastic funny business in our world, and horizontal multiplication works only when it works.

6. Need for more

In this essay we have motivated the elements of PrBST theory, and we have applied it to the Clock-Pluto problem with which we began. Even if PrBST theory can treat only problems as simple as that one, it remains worthwhile, since it says *something* where other accounts of probability say *nothing*. In fact, however, there is much more that should be developed in PrBST theory. In particular, one has to look at much more complicated examples in order to see the role of the causal ideas of BST⁹² theory in connection with probabilities. Here are two

²⁶ In discussing the universality of the Markov property, it is necessary to keep in mind that we are not speaking of a “causal ordering” of “variables” in some imperfectly defined sense of these words. The “causal order” between the “variable” *smoking* and the “variable” *cancer* is of course real and of great importance, but not easily made a matter of fully rigorous and objective definition. In contrast, the causal ordering of basic transitions is, in the spirit of Frege, sharply defined in terms of the indeterministic causal ordering of point events as described by BST⁹² theory.

small examples. (1) (Weiner, 1997) explores the following case. Let I be upper bounded, and let $Sups(I)$ be the set of all its historical suprema according to BST^{NoMFB} . How should we calculate the probability distribution on $Sups(I)$, given I ? Do the probabilities add up to 1? (2) Let it be that E can be intuitively considered both an initial event and an outcome event. Suppose we are given $pr(\mathbf{I} \mapsto E)$ and $pr(E \mapsto \mathbf{O})$. Under what conditions can we reasonably calculate $pr(\mathbf{I} \mapsto \mathbf{O})$? (The answer “always” is false.) More generally, one wants to consider various principles of probability when put in explicitly causal form. Because of length considerations, however, we beg leave to forbear pursuing this and other explorations.

7. Appendix

We recall some BST^{92} concepts and notation.

Fundamental. Our World (our only world) is represented as a set of point events together with the indeterministic and relativistic causal ordering $<$ on these point events. For notation we use e or p for a point event, h for a history (a maximal upward-directed set of point events, not to be confused with a “world”), $Hist$ for the set of all histories, H for a proposition (set of histories), with consistency and contingency thereof defined, \mathbf{H} as a set of propositions with joint consistency thereof defined. Also defined are propositions as (historically) necessary and sufficient conditions (CC §3.1).

Types of events. In thinking about indeterminism, one needs carefully to consider and distinguish initial events, outcome events, and transition events. These words are jargon, but the ideas are essential.

Initial events. An *initial event* is a set of point events all of which belong to some one history. We use \mathbf{I} for an initial event and I for the special case of a non-empty upper-bounded chain of point events. We may identify $\{e\} = e$ (CC §3.3).

Outcome events. Detailed causal theory requires complexity here. O^* ranges over three fundamentally different kinds of outcome events. (1) We use O for an *outcome chain*, which is a nonempty and lower-bounded chain of point events. Here too we may identify the unit set of a point event with the point event itself, but in the “outcome” case we conventionally use the letter p instead of e . (2) \mathbf{O} is a *scattered outcome event*, which is defined as a nonempty set of outcome chains that are consistent in the sense that they can all begin together: There is a h such that $h \cap O \neq \emptyset$ for every $O \in \mathbf{O}$. We may identify $\{O\} = O$. Scattered outcome events \mathbf{O} , which may be scattered either space-like or time-like, are often said to be *non-disjunctive*. (3) \mathbf{O} is a *disjunctive out-*

come event, which is defined as a nonempty set of pairwise-incompatible scattered outcome events (one may exclude a unit set as trivial). For outcome events, see CC §3.4. For an account of the importance of the distinction between non-disjunctive and disjunctive outcome events (not mere propositions), see the discussion in CC of its Figure 1.

Basic outcomes. Of special interest are some “basic” notions. O is a *basic (chain) outcome of e* \leftrightarrow_{df} e is a proper infimum of O . \mathbf{O} is a *basic (scattered) outcome of e* \leftrightarrow_{df} for some history h containing e , \mathbf{O} is the set of all basic chain outcomes overlapping h . $\Omega_e \leftrightarrow_{df}$ the set of all basic scattered outcomes of e , a partition of all outcome chains such that $\inf(O) = e$.²⁷

Transition events. In BST⁹² theory, a transition event $\mathbf{I} \mapsto O^*$ is *just* the ordered pair of \mathbf{I} and O^* , with \mathbf{I} “appropriately” prior to O^* as defined in CC §3.5. The definition is too complicated to warrant rehearsal here. A *basic transition* has the form $e \mapsto \mathbf{O}$ for $\mathbf{O} \in \Omega_e$. When O is a basic chain outcome of e , we also feel free to save henscratches by letting $e \mapsto O =_{df} e \mapsto \Omega_e \langle O \rangle$, the point being that in BST⁹² it is provable that the occurrence of O is the same thing as the occurrence of $\Omega_e \langle O \rangle$.

Occurrence propositions. These are separately defined for each type of event, initial, outcome, or transition in various subsections of CC §3. Intuitively to say that an initial-type event \mathbf{I} occurs in a history h is to say that it comes to completion in h , whereas to say of an outcome-type event O^* that it occurs in h is to say that it begins to be in h . (It is provable that for every transition event $\mathbf{I} \mapsto O^*$, the occurrence of the initial is a necessary condition of the occurrence of the outcome.) The intuitive reading of the truth-in- h of the occurrence-proposition attached to the entire transition event $\mathbf{I} \mapsto O^*$ is that *if \mathbf{I} occurs in h then O^* occurs in h* , with the if-then given its truth-functional reading. A *basic (propositional) outcome of e* is the occurrence proposition for some member of Ω_e , and Π_e may be defined as the set of all such.

References

- Bell, J. S.: 1964, ‘On the Einstein-Podolsky-Rosen paradox’. *Physics* **1**, 195–200.
 Belnap, N.: 1992, ‘Branching space-time’. *Synthese* **92**, 385–434. See (Belnap, 2003a).
 Belnap, N.: 2002a, ‘Double time references: Speech-act reports as modalities in an indeterminist setting’. In: F. Wolter, H. Wansing, M. de Rijke, and M. Zakharyashev (eds.): *Advances in modal logic*, Vol. 3. Singapore: World Scientific

²⁷ CC speaks of “basic primary outcomes,” verbiage useful when working around modal funny business; but here, where there is none, we may drop “primary.”

- Co. Pte. Ltd., pp. 37–58. A preprint of this essay may be obtained from <http://www.pitt.edu/~belnap>.
- Belnap, N.: 2002b, ‘EPR-like ‘funny business’ in the theory of branching space-times’. In (Placek and Butterfield, 2002), pp. 293–315. A preprint of this essay may be obtained from <http://philsci-archive.pitt.edu>.
- Belnap, N.: 2002c, ‘A theory of causation: *Causae causantes* (originating causes) as inus conditions in branching space-times’. A preprint of this essay may be obtained from <http://philsci.archive.pitt.edu>.
- Belnap, N.: 2003a, ‘Branching space-time, postprint, January, 2003’. This is a postprint of (Belnap, 1992) that includes a number of additional explanations and a little re-structuring. It may be obtained from <http://philsci-archive.pitt.edu>.
- Belnap, N.: 2003b, ‘No-common-cause EPR-like funny business in branching space-times’. *Philosophical studies* **114**, 199–221. A non-quotable preprint may be obtained from <http://philsci-archive.pitt.edu>.
- Belnap, N., M. Perloff, and M. Xu: 2001, *Facing the future: Agents and choices in our indeterminist world*. Oxford: Oxford University Press.
- Coffa, J. A.: 1973, ‘Foundations of inductive explanation’. Ph.d thesis, University of Pittsburgh.
- Eagle, A.: 2003, ‘20 arguments against propensity analyses of probability’. Draft of 4 July 2003. It may be obtained from <http://philsci-archive.pitt.edu>.
- Fleming, G. N.: 1965, ‘Covariant position operators, spin, and locality’. *Physical review B* **137**, 188–197.
- Müller, T.: 2002, ‘Branching space-time, modal logic and the counterfactual conditional’. In (Placek and Butterfield, 2002), pp. 273–291.
- Müller, T.: 2003, ‘Probabilities in branching space-times’. (Tentative title).
- Oksanen, M.: 2003, ‘Causality and events in branching space-time’. Unpublished.
- Placek, T.: 2002, ‘Partial indeterminism is enough: a branching analysis of Bell-type inequalities’. In (Placek and Butterfield, 2002), pp. 317–342.
- Placek, T.: 2003a, ‘On Belnap’s branching space-times’. Unpublished. Department of Philosophy, Jagiellonian University.
- Placek, T.: 2003b, ‘On objective transition probabilities’. Working document.
- Placek, T. and J. Butterfield (eds.): 2002, *Non-locality and modality*. Dordrecht: Kluwer Academic Publishers.
- Salmon, W. C.: 1984, *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press.
- Sharlow, M. F.: 1998, ‘A new non-Hausdorff spacetime model for resolution of the time travel paradoxes’. *Annals of physics* **263**, 179–197.
- Sharlow, M. F.: 2003, ‘The quantum mechanical path integral: toward a realistic interpretation’. Unpublished and copyright. See <http://www.eskimo.com/~msharlow>.
- Thomason, R. H.: 1970, ‘Indeterminist time and truth-value gaps’. *Theoria* **36**, 264–281.
- van Fraassen, B. C.: 1980, *The scientific image*. Oxford: Oxford University Press.
- Weiner, M.: 1997, ‘Probability in branching space-time’. Unpublished. Department of Philosophy, University of Pittsburgh.
- Weiner, M.: 2003, ‘Testimony: Evidence and responsibility’. Ph.d. thesis, University of Pittsburgh.
- Xu, M.: 1997, ‘Causation in branching time (I): transitions, events and causes’. *Synthese* **112**, 137–192.