

1

Introduction

The subject of this book, to put it broadly, is the analysis of *real possibilities*, or *indeterministic possibilities for the future*. The theory of Branching Space-Times that we develop here describes with mathematical rigor how real possibilities can be anchored in a spatio-temporal world that is rudimentarily relativistic. The picture we attempt to paint represents indeterministic events as happening locally in our spatio-temporal world, and it portrays indeterminism modally, via *alternative possibilities for the future* that are open in particular circumstances. This idea can be traced back to Aristotle. In *De Interpretatione* 9 (19^a 12–14) he says of a particular cloak, under particular circumstances, that the cloak might wear out, but that it could also be cut up first, that is, before wearing out. The currently dominant analysis of modalities in terms of possible worlds will represent the possibilities open to the cloak by multiple possible worlds, in some of which the cloak is cut up, whereas it wears out in the others. That analysis thus (1) postulates many possible worlds, (2) includes the notion of an actual world, and (3) identifies the actual world with one of the possible worlds. Branching Space-Times opposes this representation. Given that possible-worlds theories assume these three claims, Branching Space-Times is thus not a possible-worlds theory.

It is true that, in general, talk of possibilities only makes sense before a contrast between possibility and actuality. This idea can also be traced back to Aristotle, who observes in *Metaphysics* 9 (1048^b 4–6) that, together, the actual and the potential form an antithesis. However, it is not generally true that actuality has to be structurally like the possibilities with which it is contrasted, that alternate possibilities have to be alternatives *to* a given actuality. In the case of real possibilities, actuality is given as the concrete situation here and now, while the possibilities are alternatives *for* the future of that concrete situation, with none of these possibilities being actual (yet). What is actual (i.e., the concrete given situation) is structurally different from the alternative possibilities for the future.

1.1 Real possibilities

Some things are really possible, while others are not. Real possibilities are special because they are future-directed. Our talk of possibilities does not always distinguish them from other types of possibilities, but the distinction is philosophically important and easy enough to make. Let us start with an example.

Consider Alice, who is sitting in a restaurant in Pittsburgh talking to the waiter who is taking her order for lunch. At the moment we are considering, the waiter has asked, “Do you want fries with that?”, and Alice is thereby prompted to give the answer “yes” or “no”. It is really possible for her to give either answer. At these particular circumstances (but not later), she has two alternative possibilities for the future, saying “yes” and saying “no”. Many other things are, however, not really possible. For example, it is not really possible for her to turn her glass of water into a plate of fries. Nor is it really possible for her to have dinner in Kraków later on the same day, because there are no means of transportation available that could get her there in time (note that at noon in Pittsburgh it is already 6 pm in Kraków). If Alice was in Konstanz, the situation would, of course, be different: given a concrete situation of having lunch in Konstanz, it is really possible to have dinner in Kraków on the same day.¹

Real possibilities as alternatives *for* rather than alternatives *to*. Put abstractly, given the situation at hand at some concrete moment, some things can really follow from those circumstances, while others cannot. Real possibilities are always future-directed and tied to a concrete situation constituting the “here and now”. Real possibilities are, to make a point about English usage, alternatives *for* the future, for what can happen in the future. It makes no sense to say that they are alternatives *to* the future, because the future, being in the future, has not yet happened. None of the possibilities for the future is actual yet.²

¹ From Pittsburgh’s Shadyside, it would take Alice at least an hour (more realistically, two to three) to be airborne, and the distance is such that current commercial aircraft take at least 9 hours to complete the journey to Kraków Airport. (More realistically, unless Alice has access to a private jet, she will have to stop over at least once, which would prolong the journey still further). From Konstanz, on the other hand, a direct flight from Zürich takes just 2 hours, and there is no difference in time zone.

² This point is made forcefully in Rumberg (2016b).

Real possibilities are special. We talk of other types of possibilities as well, and for many of those, alternative possibilities *are* alternatives *to* a given actuality. Logical, metaphysical and natural (physical) possibilities all provide examples in which actuality is structurally like the alternative possibilities. We provide an illustration in terms of epistemic possibilities, as these form a crucially important class.³ Let us thus consider Bob, who is looking for his keys. For all he knows, they could be in his office, or they could be in his car. It is perfectly adequate to say that it is possible that the keys are in Bob's office, and it is possible that they are in his car. In actuality, they are either in one place or the other, and that is a settled matter. In fact, Bob operates under the explicit assumption that one of the options is settled to be true, and the other is settled to be false; he just does not know which is which. Bob is considering epistemic possibilities rather than real possibilities. What is epistemically possible for him is what is compatible with his evidence, and that will change as he learns things. Bob faces an actual situation and a number of alternatives to that actual situation, which he cannot distinguish based on his available evidence. He thus faces the task of finding out which of these alternatives is the actual one. In actuality, it is already settled that the keys are either in his office or in his car, and let us hope for Bob's sake that he will find out quickly which it is.

From an abstract point of view, Bob's predicament is one of self-location. Given his evidence, there is a space of possibilities that includes actuality.⁴ His task is to rule out the non-actual alternatives so that he knows his location in that space of possibilities, which will enable him to retrieve his keys. The way he does this is by looking, by acquiring new evidence that will narrow down the space of alternative possibilities. A valid formal approach to modeling Bob's predicament is in terms of possible worlds (or perhaps,

³ Linguistic data in fact seem to support the thesis that generally, when we say "it is possible that...", we are pointing to epistemic and not to real possibilities (Vetter, 2015). Real possibilities are more often expressed via the so-called root modality, as in "this dog could bite" or "Alice can choose to have fries". The root modality generally expresses either real possibilities or deontic possibilities, i.e., what one is allowed to do. It may be interesting to note that actuality need not be among the scenarios that are deontically possible (i.e., admissible) either.

⁴ Generally speaking, epistemic (or rather, doxastic) possibilities might *exclude* actuality: Bob might have mistaken beliefs, which he would need to revise in order to properly self-locate. Bob's evidence might after all be false: the keys might in fact have been taken by a prankster in his office. We rule out these complications, noting that there is a large body of literature on belief revision that deals with cases in which agents have to revise their evidence (see, e.g., Spohn, 2012). For our purposes, the important point is that even if these issues are taken into account, it remains adequate to represent that relevant space of possibilities via possible worlds.

smaller entities such as possible situations). There is Bob's actual situation, and there are alternatives to it; actuality is just one of the possibilities.⁵

1.2 Representing possibilities via branching vs. possible worlds

Many kinds of possibilities can be represented via possible worlds. But for real possibilities, a representation in terms of possible worlds distorts the picture. Consider Alice in the restaurant again. Actually, Alice has yet not answered the waiter, so none of the alternatives for the future that she is facing has been realized yet. She can really give either answer. Her task is to decide what to do, whether to order fries or not. Her actual situation does not yet include a decision; her task is not to find out what is *true*, but to find out what is *good* for her, and then to make it happen. If one tries to formally represent Alice's possibilities for the future via possible worlds, there has to be an actual world among them, which means that one complete temporal course of events has to be actual. This, however, distorts the structure of temporal actuality, since none of the possibilities for the future is yet actual. A more adequate formal representation is not in terms of possible worlds, but in terms of a branching structure of possible courses of events. One such course of events represents Alice's answering "yes", and another one of those courses of events represents Alice's answering "no". At the branch point, which represents Alice's actual situation, both courses of events are really possible as continuations of the actual situation, and none of these possibilities has yet been actualized.

Branching Space-Times builds on the idea of multiple alternative possibilities open in concrete situations. It identifies *Our World* with the set of all real possibilities accessible from a given concrete event. The set is constructed in indexical fashion by starting with a concrete actual event and then including all of the really possible events that are accessible from it. For instance, taking Alice's situation in the restaurant in Pittsburgh as actual, we can say in indexical language that it is really possible that she will answer yes, or that she will take a walk after lunch, or that she will come back to the

⁵ This verdict does not change if one explicitly includes time in the picture, as is common, for example, for tasks involving temporal self-location (see, e.g., Spohn, 2017). In that case, both actuality and the alternatives are properly represented via so-called centered possible worlds, which are possible worlds in which one point in time is singled out as the current one. The formal machinery is more or less the same.

same restaurant in a year's time. We can also say that it was possible that she would not go to the restaurant, but rather skip lunch, or that she could have left Pittsburgh for Kraków a week ago, or that at her first day at school many years back, she could have put on her blue socks instead of the green ones she actually put on then. Based on the plethora of real possibilities open at various really possible events accessible from Alice's concrete situation, we thus arrive at the concept of a complete possible course of events or, as we will say, a *history*. The farther back we go, the larger the set of possible courses of events that becomes accessible. Since courses of events are differentiated by alternative possibilities open in concrete situations, any two courses of events overlap, sharing a common past. A formal structure representing alternative real possibilities thus forms a unified whole, not a set of separate alternatives. It is possible to map out the whole structure of real possibilities by starting from one really possible point and going back and forth along the branching histories. The unity of the structure can be described in indexical terms starting from *any* one point.

It is illuminating to see what happens if one tries to force a formal representation of real possibilities in terms of possible worlds. Possible worlds are themselves complete, separate courses of events. It is, therefore, easy to extract a possible worlds representation from a branching representation; just put all the possible histories side by side without any overlap. Actuality then has to be represented by one of these worlds, perhaps together with an indication of a temporal now. How is Alice represented? Without entering the somewhat controversial topic of representing individuals,⁶ the sensible majority option seems to be that Alice, our Alice that we can point to, is an inhabitant of precisely one of those possible worlds. Other worlds may contain Alice-lookalikes ("counterparts"), but not her. In this picture, Alice's decision of what to do is represented in exactly the same way in which we represent Bob's looking for his keys: there is an actual world, Alice's actuality, for which the past as well as the future are both fully fixed. Alice is just uncertain about her location in the total space of possibilities. Actuality is guaranteed to be one of them. Her actual past, which is (or at least could be) known to her, is compatible with many other possible worlds, while her future is such that she just has no information about it. In that possible worlds rendering of the situation, deciding what to do—which is tied, some way or other, to finding out what is *good*—amounts to gathering more evidence—that is,

⁶ See Lewis (1986a, p. 199) and Kripke (1980, p. 45).

finding out what is *true*. Being tied to one possible world, Alice cannot, as it were, jump worlds to arrive at a better one. There is no agency, no room for practical rationality; there is just theoretical rationality tied to finding out what is true. This picture does not capture what is special about agency and therefore appears inadequate. A branching framework, on the other hand, can picture agents choosing between different possibilities for their future and therefore leaves room for them to select the good over the bad.

What we have just claimed might sound contentious, but it suffices, we think, to at least motivate the quest for a detailed representation of real possibilities in a branching framework.

1.3 Some thoughts on our modally rich world

We turn next to the philosophical vision underlying our construction of branching structures.

Our World. *Our World* is very big. For one, many things have already happened in *Our World*, and there is more to that particular past of *Our World* than these happenings. Even if we had a full specification of all of them, that would not provide us with complete information about the past. Things that once were really possible but did not happen belong to a full specification of *Our World* as well. Such real but never actualized possibilities are not necessarily derivable from what has happened. Similarly, all things that can really happen later on also contribute to a full characterization of *Our World*. *Our World* therefore has modal aspects that are not reducible to its non-modal features. Branching Space-Times decides to describe our modally rich world in terms of really possible events. Such events find their place in a coherent formal structure representing alternative real possibilities. That structure forms a unified whole, containing all really possible courses of events. And any two such courses of events overlap, sharing a common past.

Before we describe the formal structure in detail in subsequent chapters, we present a vision that underwrites the concept of a modally rich world, and we reflect on its coherence in particular. That vision is not strictly necessary to understand and to follow our formal constructions and their applications in the chapters to come. We nevertheless lay it down here to specify a useful metaphysics for the theory of Branching Space-Times.

To repeat, to arrive at the totality of all really possible events, one needs to start with some actual happening, like your actual eye-blinking (assuming you just blinked). There are events that can really happen later (i.e., after your eye blinking); we say that they belong to the future of possibilities of this actual event. There are also events that have already occurred, and events that could really have occurred. Your actual eye-blinking is linked to other events that are, from its perspective, really possible, via a pre-causal relation that can be expressed by a phrase like “something can really happen after something else”. This relation is used to capture the totality of all really possible events from the perspective of our actual event. What we have just said might suggest that, in order to account for the totality of all really possible events from the perspective of a given actual event, it is enough to shift the perspective back in time and then forward along one of the open possibilities. It will transpire, however, that in a theory of branching histories that represents both space and time, we will need not just a *V*-shaped indexical reference back and forth, but an *M*-shaped zig-zagging chain of indexical references—see Fact 2.4. We therefore take *Our World* \mathcal{W} to be the totality of possible events that are linked to an actual event by any *M* zig-zagging curve admitted by the pre-causal relation.⁷

The word “our” above is not merely a stylistic ornament: it reflects the construction and the fact that we have to start with the actual here-and-now. That might raise the concern that there could be many different *Our Worlds*, depending on which actual event we start with. This option is, however, excluded by our requirement that all possible courses of events are linked via a common past. The totality of possible events is therefore not relativized to an actual event. In other words, the unified structure of *Our World* can be described in indexical terms starting from *any* really possible event.⁸

⁷ This is reminiscent of David Lewis’s (1986a, p. 208) way of demarcating a possible world: its elements “stand in suitable external relations, preferably spatiotemporal”. Note, however, that while Lewis’s demarcation is non-modal, our characterization is modal, as indicated by the phrase “something *can really happen* in the future of something else”. Furthermore, as our characterization is not spatio-temporal, BST is not threatened by current developments in physics that take space-time to be an emergent phenomenon. For weird consequences of demarcating a possible world in Lewis’s way, see Wüthrich (2020).

⁸ Of course, even given the uniqueness of *Our World*, agents can be *epistemically* uncertain about what is actual and what is or was really possible. Such uncertainty could be modeled via a *set* of branching models representing different epistemic alternatives. The uniqueness of *Our World* then translates into the fact that only one from among the set of these models can contain any actual and any really possible events. In this book, we are not concerned with epistemic possibilities, and so we leave the issue for another occasion.

Events. We set out our vision of *Our World* thus far in terms of (token) events without further elucidation. We are aware that talk of events is ambiguous: there are large events as well as small events, and some events have a modal multiplicity in that they can occur in different ways (e.g., faster or more slowly), whereas other events can only occur in one way. In what follows, we take idealized point-like events as the basic building blocks of our construction, in analogy to Euclidean points or to the physicists' mass points. We assume that our point-like events cannot occur in different ways, that is, they do not have the mentioned modal multiplicity. Starting from point-like token events, we then construct other varieties of events, including so-called disjunctive events, which *can* occur in different ways.

An event, as we understand it, is not a place-time or a collection of such; it is a happening. It has a time and a place, which partly describes its locus in our world—but that is not enough to confer uniqueness. An event has a concrete past and a concrete future of possibilities. Nor is the past of an event a mere array of times and places: such a past also consists in concrete events. Events, as we construe them, are as concrete as your actual eye-blinking, they have a definite relationship to this very eye-blinking and are related to all other point events that from our point of here-and-now are either actual (such as those in your past) or really possible. A consequence that we want to draw from this characterization is that token-level events cannot be repeated. In contrast, a type-event is repeatable.⁹ To make one final point about events: despite the basic role they play in our theorizing, the Branching Space-Times project should not be read as a commitment to a reductionist program in ontology that aims to reduce everything to events. Our decision to focus on events is driven by our diagnosis as to which objects are known to be amenable to formal analysis: we know how to formally analyze events, but we also know that adding processes, or enduring objects, makes the task much harder. So events is what we start with: we begin with possible events, idealized as point-like, and then construct other varieties of events as well as some other kinds of objects.

Histories. The totality of all really possible events contains compatible events (i.e., events that can occur together) and incompatible events (i.e., events that cannot occur together). Of particular interest are incompatible

⁹ In some sense, the non-repeatability of token events might still be contentious, as there might be closed causal loops—see Section 9.3.6 for discussion.

events that are local alternatives for the future, like Alice's alternative possible responses "yes" and "no" to the question about her lunch. Which events are compatible? We follow what we call the later witness intuition, according to which two events are compatible only if they belong to the past of some really possible event. Based on the criterion provided by this intuition, we then define histories as maximal sets of compatible events, where "history" is our technical notion for the informal "course of events". We consider histories a useful device for tracking local compatibilities and incompatibilities, but we tend to minimize their ontological significance. That is, we take it that local alternatives occurring in space and time are ontologically important, whereas histories merely offer a useful way of conceptualizing alternative possibilities. There are less demanding conceptualizations that better serve the local vision underlying Branching Space-Times. Two such concepts, alternative transitions (Müller, 2014; Rumberg, 2016a,b) and alternative possible continuations (Placek, 2011), have already been proposed for the branching framework. We nevertheless work here in terms of histories, mostly for reasons of simplicity. Note also that given an arbitrary criterion of (local) compatibility, one can typically define maximal sets of compatible events,¹⁰ which means that the formal basis for defining histories is available anyway.

As we have already emphasized here, courses of events overlap, and indeed our postulates require any two histories to have a non-empty intersection. More precisely, if a structure has more than one history, it contains an object (a point event or a chain of points events) at which histories branch. But what is it, precisely, that branches? There are two ways of thinking about branching that should be avoided: (1) individual space-times branch or (2) *actual* courses of events branch. As for mistake (1), since histories are identified with concrete possible courses of events, there is more to a history than its spatio-temporal structure. Branching Space-Times in fact allows that all the different histories (all the concrete individual space-times) have the same space-time structure. For an example, see the so-called Minkowskian Branching Structures discussed in Chapter 9.1. Clearly, individual space-times do not branch in structures of this kind, as every history has the same

¹⁰ Set-theoretically, the Axiom of Choice is required. In this book we freely use the Axiom of Choice and equivalent principles such as the Zorn-Kuratowski lemma or Hausdorff's maximal principle. For some details, see note 8 on p. 30.

non-branching, Minkowskian space-time structure.¹¹ It is spatio-temporal *histories* that branch. To address the second mistake, a point about language might help. While explaining the pre-causal relation, we used the phrase ‘what can really happen later’, not ‘what will happen’. To illustrate, focusing on an actual event of coin tossing, the coin can really land heads up and it can really land tails up. Thus, it is alternative *possibilities* open in particular circumstances, or more generally, alternative *possible* courses of events, that branch.

Temporal directedness. We insist on the significance of alternative possibilities for the *future* in our branching framework. We assume that the temporal directedness of the pre-causal relation is objective and derivable from the modal features of our world, viz., from the distinction between a settled past and an open future. In particular, if there is no room for real possibilities in *Our World*, there is no directedness, and nothing can distinguish the past and the future of any event. In the presence of real possibilities, however, the future is modally distinguished from the past.

Given that there are real possibilities, we need to take a stance on how to represent the world-lines of point-like objects in a branching structure. Such a world-line (a trajectory) is defined in Branching Space-Times as a maximal chain of events, where a set of events is called a *chain* if any two elements of it are comparable by the pre-causal relation. With real possibilities present, some pairs of trajectories bifurcate, forming a *Y*-shaped figure. Two such bifurcating trajectories share a common “trunk” that is adjacent to two disjoint “arms”. Given continuity (in line with the standard assumptions in physics), either (i) the trunk has a maximal element and each arm has no minimal element, or (ii) the trunk has no maximal element, but each arm has a minimal element. It turns out that both options are formally viable. For our theory, we decide the issue globally, by assuming different postulates governing lower bounded chains and upper bounded chains. A lower bounded chain has to have a (unique) greatest lower bound, whereas an upper bounded chain might have multiple minimal upper bounds, as in option (ii) for bifurcating trajectories. As a result of these diverging postulates, if we start with a branching structure and then systematically flip the direction, typically the resulting object will not be a branching

¹¹ For the record, Minkowskian Branching Structures are not the only option. Branching Space-Times also allows that different histories have vastly different space-time structures (see, e.g., the construction suggested in Exercise 2.5).

structure—it will violate the infima and suprema postulates for bounded chains. For a further discussion of temporal directedness, see Chapter 2.4, and for further discussions relating to the philosophy of time, see Chapter 10.

1.4 Branching in the landscape of possible-worlds theories: Some comments on modal metaphysics

Branching Space-Times, as we said, is not a possible-worlds theory in the standard sense: it does not posit a multiplicity of possible worlds, and thus it does not single out an actual world from among that multiplicity. The main reason for that difference is that Branching Space-Times is meant to describe a kind of possibility that is different from the usual target of possible-worlds analyses. We have argued that possible-worlds frameworks do not do justice to real possibilities, and we said that the different targets of the two approaches give rise to deep formal differences of how the two approaches represent modalities (see Section 1.2). Yet, it is useful to describe the branching approach within the larger landscape of positions in modal metaphysics. One might think that this should be straightforward. After all, the positions are differentiated by their stance with respect to just a few issues, and so, by learning the branching theorists' response to them, one should be able to locate the theory on the map of standpoints in modal metaphysics. The most important issues are the following: Is the theory intended to be reductive with respect to modalities, by reductively explaining them in non-modal terms of some sort? Are possible worlds (or histories, or any objects standing for full possible courses of events) thought of as actual, and, if not, how is the distinction between actuality and mere possibility explained? The next big question is how individuals are represented in these possible worlds or histories, and especially how the modal features of agents that pertain to exercising their agency are described. Finally, there is a question concerning the status of the laws of nature vis-à-vis possibilities: Are the possibilities dictated by the laws of nature and particular circumstances, or is it precisely the other way round, so that (real) possibilities delineate what the laws are in *Our World*?

It is true that the branching approach and the possible worlds approach do differ, more or less strongly, with respect to the mentioned issues. But the assimilation of branching with possible-worlds theory is also somewhat treacherous, and for two reasons. First, in the interest of maximal formal

rigor, Branching Space-Times theory is very frugal with respect to its primitive notions. Having only a few primitive notions increases the transparency of our formal constructions. More advanced concepts are added later on, and they need to be constructed in terms of primitive objects. Now it is exactly such advanced concepts that are needed to draw comparisons between branching and possible-worlds theories. But these advanced concepts, being derived rather than primitive, are not well-suited for drawing fundamental distinctions between Branching Space-Times and possible-worlds theories.

The second and more important reason is that branching and possible-worlds theories respond to different data. They have somewhat diverging aims and different criteria of success. The theory of Branching Space-Times belongs to metaphysics. It picks out the notion of alternative possibilities for the future as its starting point and assumes that this notion is clear enough to permit a non-controversial formalization. It then uses that notion to mathematically analyze local indeterminism occurring in relativistic space-time. The hope of the branching theorist is that the theory's mathematical elegance, its broad scope, the richness of its consequences, and its applicability to the analysis of problems in metaphysics and in the foundations of science will count in its favor. These virtues should thus provide a good defense of the metaphysical position that the theory formalizes. But, crucially, accounting for ways we use modal fragments of our vernacular languages is not at the top of the list of priorities of Branching Space-Times theory. It is, for example, not the theory's business to account for our practice of using counterfactuals (i.e., to account for the linguistic fact that we intuitively take some counterfactuals to be true and others to be false). In a similar vein, our linguistic practice of using alethic modalities are not the data that Branching Space-Times theory responds to. To put our cards on the table, we share the linguistic intuition that Elizabeth II might have never become the Queen of England, but that she could not have had parents different from the parents she actually had.¹² But, to repeat, such linguistic facts are not the evidential data that Branching Space-Times is meant to accommodate. This relatively low priority of linguistic data extends to the way people speak about future possibilities and actuality. Even if it turns out that our ways of speaking favor a vision with a distinguished actual future course of events (contradicting our egalitarian concept of alternative possibilities for

¹² We emphasize that these are *intuitions*. With respect to the first claim, we believe that it is an objective fact whether there really were chancy events such that, if they had happened, Princess Elizabeth would not have been crowned Queen of England.

the future), we persevere in the construction of our metaphysical theory, investigating what are the consequences of this egalitarian concept (given the assumed postulates).

The relatively low priority that Branching Space-Times theory gives to linguistic data stands in stark contrast to possible-worlds theories. These theories have either emerged from semantical theories for languages with modal operators or aim at providing such a semantics. Our modal talk is likely the most important datum that influences how these theories frame their basic metaphysical concepts, such as the similarity of possible worlds, an accessibility relation between worlds, or possible worlds themselves. For an illustration, here is David Lewis explaining how our (intuitive) knowledge of counterfactuals determines which similarity relation between possible worlds is adequate. (“Analysis 2” is his possible-worlds-based analysis of counterfactuals, for which a notion of similarity between worlds is crucial.)

[W]e must use what we know about the truth and falsity of counterfactuals to see if we can find some sort of similarity relation—not necessarily the first one that springs to mind—that combines with Analysis 2 to yield the proper truth conditions. It is this combination that can be tested against our knowledge of counterfactuals, not Analysis 2 by itself. In looking for a combination that will stand up to the test, we must use what we know about counterfactuals to find out about the appropriate similarity relation—not the other way around. (Lewis, 1986b, p. 43)

The same methodology, with a dominant role of linguistic data concerning modalities, is operative in the works of two other founding fathers of modal metaphysics, Saul Kripke and Robert Stalnaker.

Our habits of speaking might be biased toward determinism for various reasons. A possible-worlds theory could therefore be adequate in accounting for these habits, while being rather more off the mark with respect to metaphysical issues such as indeterminism. In short, what is central for a branching approach (i.e., an exercise in metaphysics to provide an analysis of local indeterminism), may be of marginal importance for possible-worlds theories. And, vice versa, the semantical enterprise, so dear to possible worlds theorists, has only limited, secondary significance for Branching Space-Times.

The semantical enterprise is also problematic because it typically brings with it a way of thinking about modalities that is foreign to real possibilities.

A particularly useful idea in modal semantics is that of an accessibility relation used to discern modalities of different grades, like logical possibilities, metaphysical possibilities, natural (aka physical) possibilities, technological possibilities, and so on. These distinct possibilities are thought of as differentiated by distinct accessibility relations, which in turn are explained in terms of different kinds of laws: what is permitted by (or compatible with) laws of a given kind, is possible in the sense related to that kind.¹³ Accordingly, we arrive at the familiar picture of increasingly smaller spaces of possibilities. There is the largest space of logically possible propositions—all those that are consistent with the laws of logic. There is a smaller space of metaphysically possible propositions, all of which are consistent with the laws of metaphysics. In a similar fashion, physical possibilities, technological possibilities, and others complete the picture. It is debatable whether this image of nesting possibilities is adequate even before bringing in the issue of real possibilities (see Fine, 2005). But our concern is real possibilities, and we are skeptical that this approach, by giving priority to laws and by characterizing ever stricter possibilities by increasingly more demanding laws, can capture real possibilities. Assume that we take laws of a certain kind as fundamental for bringing modalities into our world: Are there laws that single out real possibilities precisely? To describe a real possibility, we need to refer to a particular concrete circumstance, a particular moment in time, and a particular location in space. Can real possibilities be derived from a net of laws taken together with some initial conditions? How rich would a net of laws have to be for this to be viable? Consider Bálaž who, given all the circumstances obtaining here-and-now, could really run to the main station to catch the last train to Konstanz today. If he starts a minute later, he won't make it; under slightly changed circumstances he won't make it either. Was this feat of his really possible at the mentioned circumstance because it is compatible with some set of laws, taken together with some initial conditions? Our answer is that it was possible because of the particular circumstances obtaining in the relevant region of *Our World*. Perhaps one could derive the required laws of nature from the real possibilities that obtain, distilling the laws, so to speak, from the dispositional and modal features of our world. But the net of laws that would be required to determine the real possibilities that obtain for our runner in the given concrete circumstances

¹³ In contrast, in *Branching Space-Times*, a notion of accessibility can be defined in terms of the relative location of points in a branching structure, without reference to laws; see Müller (2002).

would be immense and surely beyond our comprehension. This net of laws looks very different from the laws that we know from the sciences, or from philosophical accounts of laws. Just compare the attempt at providing a law-based account of our runner's real possibilities with a smooth account concerning physical possibility (e.g., that for a photon it is physically possible to travel from the Moon to the Earth within 3 minutes because that does not contradict the relativistic limitation on the speed of light or any other known candidates for laws of nature). Branching theory thus suggests that real possibilities do not fall in the mold of law-given modalities. The suggestion, therefore, is to take real possibility as a primitive, non-reducible notion and study it by constructing a formally rigorous theory.

So much for the differences between Branching Space-Times and possible-worlds theories. Given these differences and the frugality of Branching Space-Times, we believe that it is far from helpful to attempt to locate this theory on the map of possible-worlds theories. In order to contribute to the discussions in the literature, we nevertheless end this section with some remarks on three issues that call for our particular attention: actuality, reducibility, and the meaning of "possible worlds".

Actuality and alternatives *to* vs. alternatives *for*. Branching Space-Times theory subscribes to the semantical thesis that "actually" is an indexical word, like "here". Accordingly, if we imagine a branching-world dweller, her utterance singles out a specific piece of *Our World*, namely, the event of her particular utterance. The theory idealizes this utterance to be a point-like event. Does actuality extend any further, beyond the event of utterance? The answer is relevant for the metaphysical question of how the division of actual vs. possible is drawn in Branching Space-Times. In his account of modal realism (which is the thesis that all possible worlds are equally real), David Lewis claims that actuality somehow percolates from "me and all my surroundings" (1973, p. 86) to the whole actual world. We disagree already with this starting point. Uncontroversially, an utterance is a larger affair than a point event, but it is an innocent idealization to identify it with a point event. How can one extend actuality beyond such a point-like utterance? One idea is to extend actuality beyond the actual utterance and toward its past. Given the structure of possibilities captured by the postulates of Branching Space-Times, there are alternative possibilities for the future, but no alternative possibilities for a concrete event's past. So, since the past of the actual utterance is fixed, one can extend actuality from an actual event

of utterance downward, to include the whole past of this event. Can we go any further, taking Lewis's (1986a, p. 71) lead in appealing to spatio-temporal relations? On this proposal, if a possible event is spatio-temporally related to the actual event of utterance, then it is actual as well. We oppose this move for the simple reason that the alternative possibilities for the future of the actual utterance event are not necessarily distinguished via different spatio-temporal structures. For instance, there are specific structures of Branching Space-Times, such as the so-called Minkowskian Branching Structures developed in Chapter 9.1, in which all histories share the same space-time structure (in that case, the structure of Minkowski space-time). And yet these structures harbor different incompatible possibilities for the future of certain events. Therefore, no purely spatio-temporal relation involving the utterance can distinguish between its alternative possible futures.¹⁴

In Branching Space-Times, there is thus no actual future, and accordingly, there is no actual history. This is the expression of a basic tenet of our theory: future possibilities are alternative possibilities *for* the future and not alternative possibilities *to* an actual future. This stance contrasts with the so-called Thin Red Line doctrine, according to which there is one distinguished (actual) history and, hence, one distinguished actual future of any actual event.

One way of arguing for the metaphysics of a Thin Red Line goes via linguistic data that seem to suggest an actual future. After all, in some circumstances people utter sentences of the form "It will (actually) happen, even though it might not" (Malpass and Wawer, 2012, p. 26). One might hear in a bar, for instance, that poor Fred will actually have another beer, even though he might not. The first part of this utterance reflects on Fred's bad habits, whereas the second acknowledges an alternative following from Fred's being possibly more strong-willed. We agree that such sentences have felicitous uses, but we are skeptical of the idea that such data indicate a *metaphysical* stance about a distinguished future. It seems to us that such utterances can be accounted for in epistemic terms (e.g., by reference to the strength of expectations). We still acknowledge that at the end of the day it

¹⁴ For the record, note that structures of Branching Space-Times in which histories have different topological spatio-temporal structures do not help either (for an example of this kind of structure, see Exercise 2.5). In such a structure, the event of an utterance has to lie in a region in which the alternative spatio-temporal structures coincide. To use one of the alternative spatio-temporal structures as a criterion of actuality then betrays the very idea that motivates the picture with a topology change: both spatio-temporal structures are on a par, as each can be realized as our world develops further.

may turn out that all epistemic accounts fail, leaving us with clear evidence for people's belief in an actual future. We would take that as an indication of people's deterministic preferences. Yet, in accord with our project's assignment of a low priority to linguistic data and its avowed aim to model indeterminism, we would still not accommodate such language-based evidence for a Thin Red Line by adding a preferred history to our formal theory.

Another typical argument for Thin Red Line metaphysics rests on the desire to retain the meta-semantical intuition that, given a context of evaluation, any sentence (including a sentence about future contingents) is either true simpliciter or false simpliciter. 'Simpliciter' here means that the truth-values are not relativized to possible histories. We do not share this intuition, and we note again that it relates to linguistic or semantical matters. Their role in deciding a metaphysical issue like determinism vs. indeterminism should be fairly limited. We take it that adding a distinguished "Thin Red Line" history compromises the local indeterminism that we want to model, even if the resulting structure permits true indeterministic-looking sentences such as "it is possible that it will rain tomorrow and it is possible that it will not rain tomorrow".¹⁵

To sum up, the theory of Branching Space-Times upholds the semantical thesis of the indexical character of "actually". With regard to metaphysics, the theory holds that actuality can be ascribed to a point event (paradigmatically, an event of utterance), and, if one likes, to the past of this event. But the theory strongly opposes ascribing actuality to histories or to future segments of histories.

Eliminative analysis or modalism . An important meta-methodological issue in modal metaphysics is the following: What is the analysis of modality meant to achieve? It is typical for philosophers in the analytic tradition to deal with philosophically problematic concepts by attempting to provide an eliminative analysis. The attempted analysis aims at reducing the problematic concept to concepts that are thought to be unproblematic, or at least significantly less problematic. A well-known example is the tripartite definition of knowledge that attempts to identify knowledge with true and justified belief. Arguably, if one knows what belief is, what truth is, and what justification is, one learns from that analysis what knowledge is (if

¹⁵ Such a sentence is true at any point of evaluation because on the Thin Red Line, it either rains tomorrow, or it doesn't rain tomorrow. See, e.g., Øhrstrøm (2009) and Malpass and Wawer (2012).

one did not know it before). The controversy surrounding this example of analysis is well known. What about analyzing modality? Modal idioms might be special, in the sense that it does not seem possible to learn them by just mastering their possible-worlds analysis. Accordingly, there is the doctrine known as *modalism*, which claims that modal idioms are primitive, which implies that an eliminative analysis of modal idioms is impossible. The controversy between eliminative positions and modalism involves a number of subtleties, some related to the notions of analysis and elimination, some to drawing the line between modal and non-modal terms, and some to technical details concerning the supposed reduction. Without going into these details, we just report the consensus view that David Lewis's project of *Humean Supervenience* is intended as reductive. A non-modal analysis of laws of nature (the so-called Best System analysis) serves as its starting point.¹⁶ This is in stark contrast to Saul Kripke's (1980, p. 19) stance, which he expresses as follows: "I do not think of 'possible worlds' as providing a *reductive* analysis in any philosophically significant sense, that is, as uncovering the ultimate nature, from either an epistemological or a metaphysical point of view, of modal operators, propositions, etc., or as 'explicating' them". Kripke's stance is likely the majority view. For instance, Stalnaker's (2012, p. 30) diagnosis is that

[...] if by "analysis" one means an eliminative reduction, then I think most possible-worlds theorists (David Lewis aside) will agree with modalism, but one may still hold that possible-worlds semantics provides a genuine explanation, in some sense, of the meanings of modal expressions.

In this controversy, Branching Space-Times theory sides with the majority view, as it does not aim to eliminate modality, but rather to offer an elucidation of some modal and some non-modal notions. Starting with the primitive concept of real possibilities, the theory aims to describe local indeterminism as happening in relativistic space-time. On that basis, it aims to establish an analysis of causation in indeterministic settings and a theory of single case objective probabilities (propensities). It provides analyses of modal funny business and of non-local probabilistic correlations. These analyses are then used to address selected problems in the philosophy of

¹⁶ See, e.g., Stalnaker (2015) for an assessment.

quantum mechanics, in the philosophy of general relativity, and in the philosophy of time.

Possible worlds or alternative states of one world. The notion of *Our World* containing multiple branching histories reminds one of the controversy as to whether one should analyze possibilities in terms alternative possible worlds or rather in terms of alternative states of one actual world. When introducing possible worlds, David Lewis (1973, p. 84) writes:

I believe, and so do you, that things could have been different in countless ways. ...Ordinary language permits the paraphrase: there are many ways things could have been besides the way they actually are. ...I therefore believe in the existence of entities that might be called 'ways things could have been.' I prefer to call them 'possible worlds'.

Philosophers were quick to note that the passage from 'ways things could have been' to 'possible worlds' is far from innocuous (see, e.g., Stalnaker, 1976, and Kripke, 1980). As Kripke notes, the label 'possible world' is picturesque, but metaphorical, and potentially misleading. A more adequate terminology would be to call the entities posited by modal metaphysics "total 'ways the world might have been,' or states or histories of the entire world" (Kripke, 1980, p. 18). A possible (total) state of the world may or may not be instantiated by the actual world. Thus, 'possible ways' suggests a picture of one actual world that is capable of taking one of possibly many alternative states, whereas Lewis's phrase invokes a multiplicity of possible worlds that includes one distinguished world, the actual one. These two pictures illustrate the distinction between two varieties of views in modal metaphysics: actualism (one world with many possible total states) and modal realism (many worlds).

Returning to Branching Space-Times, its insistence on there being just one world, *Our World*, which may comprise many histories, sounds like an actualist position. However, to make any stronger claim one needs a theory of states or properties. Since Branching Space-Times, at least at its present stage of development, does not say anything about states or properties (it is purely an event-based theory), we cannot advance any stronger claim besides noting the resemblance between Branching Space-Times and actualism. The resemblance, however, is far from perfect, marred by the fact that histories are just particular subsets of *Our World*. It sounds odd to say that a part of

something, say, a part of Pittsburgh, is a state (or a property) of Pittsburgh. However, if one thinks that this case of bad English is not a major obstacle, we have no objections to understanding histories as possible ways *Our World* might be. Of course, what is then needed is a theory of states and an elaboration of histories in terms of states. The important thing which must not be lost in attempts to assimilate Branching Space-Times with actualism, is that *Our World* has a non-reducibly modal character based on multiple alternative possibilities for the future of particular events.

1.5 Outline of the book

At some point, teasing glimpses of a theory should give way to laying the groundwork for its formulation. This point has arrived and so we turn now to explaining the formal framework of Branching Space-Times.

Our book has two parts. The remaining chapters of Part I present the formal theory. Starting with Chapter 2, we introduce the Postulates of the common core of Branching Space-Times. In Chapter 3, we show that there are two options for developing this common core further, which lead to two topologically different ways for histories to branch. We introduce further defined notions in Chapter 4. The remaining three chapters of Part I introduce further formal developments of the core theory that provide the basis for applications: modal funny business (Chapter 5), causation in terms of *causae causantes* (Chapter 6), and a spatio-temporal theory of single case probabilities (Chapter 7). In Part II, we put the material of Part I to use in three concrete applications to quantum correlations (Chapter 8), to branching in (special and general) relativistic space-times (Chapter 9), and to the doctrine of presentism (Chapter 10).

1.6 Exercises to Chapter 1

Exercise 1.1. Lewis (1986a, p. 208) assumes that all elements of a possible world are to “stand in suitable external relations, preferably spatiotemporal”. Somewhat similarly, in Branching Space-Times any two point events from *Our World* are linked by appropriately combined instances of the pre-causal relation $<$ (see the M property, Fact 2.4). Discuss whether the pre-causal

relation (which is formally explained in Chapter 2.1) is a “suitable external relation” from Lewis’s perspective.

Exercise 1.2. Branching Space-Times supplies two options of how to construe actuality as a metaphysical concept: either as a token event (typically, the event of utterance), or a token event together with its past. Discuss the pros and cons of each option.