

Branching in the landscape of possibilities

Thomas Müller, Utrecht University*

Thomas.Mueller@phil.uu.nl

7 May 2011

Abstract

The metaphor of a branching tree of future possibilities has a number of important philosophical and logical uses. In this paper we trace this metaphor through some of its uses and argue that the metaphor works the same way in physics as in philosophy. We then give an overview of formal systems for branching possibilities, viz., branching time and (briefly) branching space-times. In a next step we describe a number of different notions of possibility, thereby sketching a landscape of possibilities. In the final section of the paper we look at the place of branching-based possibilities in that larger landscape of possibilities. Our main message is that far from being an outlandish metaphysical extravagancy, branching-based possibilities are epistemically as well as metaphysically basic.

Keywords: Branching time, branching space-times, modality, possibility, experiment

The metaphor of branching is useful for describing a certain type of modality: an open future can be pictured as a branching tree of concrete possibilities. In this paper we will look at the place of branching-based possibilities in the larger landscape of possibilities. The main message is that far from being an outlandish metaphysical extravagancy, branching-based possibilities are absolutely basic.

The paper is structured as follows: In §1 we trace the branching metaphor through some of its uses and argue that the uses in philosophy and in the sciences are of a piece. In §2 we briefly describe two formal frameworks for branching-based possibilities, branching time and branching space-times. §3 centers on the notion of a scientific experiment to motivate different types of possibilities, thereby giving substance to the idea of a rich landscape of possibilities. In the final §4 we argue for the basic nature of branching-based possibilities vis-à-vis other types of possibility.

*Copyright by the author. Forthcoming in *Synthese*.

1 “Branching”

The notion of branching is a forceful metaphor that has many uses; the use in logic and metaphysics is just one among many. As a first step, it is therefore appropriate to look at the literal use of the word “branching”, and to trace some of its figurative uses. In line with a general science-friendly methodology, examples will be chosen from various scientific contexts. Some relatives of “branching” such as “forking”, “ramifying” and “bifurcating”, will also be considered.

According to the *Oxford English Dictionary*, to branch is simply to bear or put forth branches, or to ramify. In its literal use, “branching” is what trees or other plants are doing as they grow. “To branch” is a verb descriptive of specific spatial processes; its noun cognate, “branch”, describes the product of such processes. Both words are also used figuratively for non-biological spatial structures such as rivers. Its use in engineering follows suit—think of roads, water supply installations, or electrical wiring. Figure 1 illustrates the engineering use. We will take all uses of “branching” descriptive of *spatial* structures or




Figure 1: From a Pittsburgh patent application for a “method of dividing and branching electric cables” (Waring, 1882).

processes, whether literal (as in biology) or figurative (as, e.g., in geology and engineering), to be unproblematic.

The common feature of these uses is that a unified spatial structure (the trunk; a piece of the river; the wires to the left in Figure 1), which has a preferred direction from some “source” to some “terminal” (from roots to tips; downstream, downhill, from source to sea; from socket to appliance), divides into many parts without losing contact with the source. The direction is crucial: we distinguish between branching—a downstream fork—and confluence—an upstream fork—, or between fission and fusion. Thus, the Ohio river doesn’t branch into the Monongahela and the Allegheny at the Point in Pittsburgh—rather, the Ohio is a confluence of the two others.

There are other, metaphorical uses of “branching” that cannot be interpreted spatially, both in the sciences and in philosophy. These uses are at least not straightforwardly unproblematic, and it is important for us to consider their

interrelation. As we are not talking about spatial structures, the question arises whether there still is a common abstract description that unites these uses. Saunders has recently claimed that the parallel use of the term “branching” in physics and in philosophy is a “coincidence in the terminology” (Saunders, 2010, 197). This claim should however be found implausible when the following examples have been considered:¹ In all these examples, we can find the features of trunk, directionality, continuity, and separation of one into many, in one form or other.

Radioactive branching. Some radioactive substances, such as actinium-226 ($Z = 89$), can decay in a number of different ways. Thus, for ${}_{89}^{226}\text{Ac}$, β^- decay into thorium-226 is the most common way of decaying (83%), but electron capture and α decay as different modes of disintegration also occur and lead to different products.

Radioactive branching was noted as a phenomenon in the 1910s, and branch ratios between different decay paths were measured; for an early review article, see Feathers (1947). Note that the metaphor of “branching” is here used independently of any spatial structure, and is directly related to indeterministic possibilities. Furthermore, branch ratios describe objective phenomena, they are not tied to any scientific theory, and they were measured well before the advent of quantum mechanics, let alone a theory of electron capture or quantum-mechanical tunneling. (Of course, quantum-mechanical modeling of the atomic nucleus helps to explain them.)

Branching stochastic processes. Kolmogorov and Dmitriev (1947) investigated what they called “branching stochastic processes”: mathematical structures that can represent temporally successive chance set-ups. The notion is retained in the subsequent literature. The authors mention that their “stochastic scheme has various applications in biology, chemistry and physics of elementary particles” (p. 5). A further important application for the resulting theorems is in controlling nuclear fission, a hot topic at the time. Nuclear chain reactions can be modeled as branching stochastic processes, and obviously there is an interest in navigating the thin line on which nuclear chain reactions neither die out nor blow up exponentially (and literally): that is how a nuclear power plant operates.

Branching in Everett’s interpretation of quantum mechanics. Everett worked on his dissertation under Wheeler in the mid-1950s. In a footnote in his famous paper, “‘Relative state’ formulation of quantum mechanics” (Everett, 1957), he illustrates his conception of branching superpositions with the image of simultaneously equally real branches, none of which would be special:²

¹No claim to historical exhaustiveness or even representativity is intended.

²Wheeler didn’t approve of Everett’s notion of branching; the note was added in proof, too late for Wheeler to toss out. Apparently Everett had used biological analogies in work

From the viewpoint of the theory, *all* elements of a superposition (all “branches”) are “actual,” none any more “real” than the rest. (Everett, 1957, 459n)

Everett argues that on his scenario, our experiences would be exactly as they are now. This interpretation of quantum mechanics has given rise to a very lively discussion, not the least about whether it should be called an interpretation at all, and on how we can understand its underlying metaphysics.

Those were some scientific examples of the use of the branching metaphor. In all these examples, the directionality is given by, or is at least parallel to, the temporal order: a unity becomes many over time. There is also a mathematical use of the branching metaphor, for which, by the nature of the subject, the temporal order can play no role. The direction here is rather given in the interpretation of the respective mathematical structures, either with a view towards (temporal or spatial) application, or perhaps via some sort of (spatial or temporal) visualization.

Branching in logic and mathematics. Many structures in mathematics and in logic are said to be branching, ramifying, or bifurcating. The AMS subject classification (AMS, 2010) currently lists 32 whole subfields of mathematics carrying these words in their title. To mention just a small number of early examples: Russell (1908) and later on, in *Principia Mathematica*, Whitehead and Russell (1910) propose ramified type theory, in which each simple type ramifies into a hierarchy of types to avoid impredicability.³ Lindenbaum and Tarski (1936) employ the notion of (non-)ramifiability⁴ of logical theories in their research on categoricity. Hopf (1942) explores the notion of bifurcation in dynamical systems. Examples can be multiplied almost indefinitely. Maybe the purest example, on which many others can be modeled, comes from the theory of orders: A partial ordering $\langle P, < \rangle$, where $<$ determines a natural “forward” direction, is said to be branching (to the right, or forward) iff there are $x, y, z \in P$ for which $x < y$ and $x < z$ while neither $y < z$ nor $z < y$ nor $y = z$. Trees—specific branching partial orders—are ubiquitous in logic and in computer science.⁵

There is a number of different *philosophical* uses of the branching metaphor: uses that point to spatial structures, but also properly temporal ones. Tree-like

leading up to his thesis, e.g., the image of an amoeba splitting in his “Probability in wave mechanics”. By the time of the thesis, however, he had dropped talk of amoebas; see Byrne (2010, 525f.).—The image of “many worlds” was introduced later, by DeWitt (1970).

³See Church (1956, § 58) for details.

⁴The German original has “Gabelbarkeit”. Thanks to Jan Woleński for drawing my attention to this example.

⁵In fact, trees are usually defined via backward *non-branching*, i.e., the non-existence of elements x, y, z for which $x < z$ and $y < z$, but x and y are incomparable. This leaves only forward branching, with linearity as a degenerate special case.

structures that can be read spatially occur frequently in philosophy, e.g., in the form of various taxonomies like the *arbor porphyrii*. We will consider exclusively the philosophical image of branching in (or of) time as a representation of possibilities, and its refinements.

Branching in (of?) time. The image of a branching future as the home of open possibilities appears to be quite natural. We subscribe to the picture that things may go one way or the other—a spatial branching metaphor. This is a well known topic in literary writing. Borges is famous for his short story, “The garden of forking paths”, written in 1941 and translated into English in 1948 (the Spanish original has “se bifurcar”). The story overlays a spatial branching structure of garden paths with the metaphorical branching of time.⁶ Incidentally, Borges admits both forward and backward branching.

Branching time As far as we know, the image of a branching future of possibilities appears in logic for the first time in 1958, in a letter that Saul Kripke, then a beginning student on his way to Harvard, wrote to Arthur Prior.⁷ In 1957, Prior had just published his Oxford *John Locke Lectures* of 1955/56 as the book *Time and Modality* (Prior, 1957), which marks the beginning of his project of tense logic. Questions of the interaction of time and modality were crucial for this project from the beginning. Indeed, Prior argues that his “tense logic [...] is a modal system” (p. 12). The basic idea at the time was that possibility should be something like truth at some time, so that ϕ is possible if and only if ϕ is either true now or will be true at some time in the future. This was called “Diodorean possibility”, after an association with the so-called Master Argument of the Megarian philosopher Diodorus Cronus. The definition of the Diodorean possibility operator $Poss_D$: is

$$Poss_D:\phi \Leftrightarrow_{df} \phi \vee F\phi,$$

where F is Prior’s future-tense operator, “it will be the case that”.⁸ The metaphor that possibility is truth somewhere or somewhen is strong in many philosophers; the image is kept up, e.g., in Lewis’s modal realism.⁹

Kripke suggested, instead of the Diodorean approach, to unite time and modality as two *separate* interpretations of a *single* relational structure,

⁶Cooper, who had independently discovered Everett’s theory together with a student (Cooper and Van Vechten, 1969), explicitly references Borges’s story, saying that in it, “the poet has preceded the scientist” (Cooper, 1976, 45).

⁷The letter will hopefully be published in a collection of writings on Prior edited by Peter Øhrstrøm.

⁸Prior actually defines the possibility operator via metric tense logic (Prior, 1957, 13), but this is immaterial for our discussion here.

⁹See, e.g., Lewis (1986).—Apparently Peter Geach was the first to suggest a space-traveling metaphor for the “accessibility” relation that was being introduced into the semantics of modal logics around 1956; he spoke (actually in a somewhat mocking tone) of a “dimension-jumping vehicle” that would take one along that relation between possible worlds; see Prior (1962).

creating what may be the first diagram of “branching time” in logic. (Kripke does not use the word “branching”, but he calls his diagram a “tree”, displaying the root at the top as is common in logic.)

Branching time was researched into in the 1960s, and by 1970, there was a useful semantic framework for a temporal-modal language based on branching time (Thomason, 1970, 1984). Branching time allows one to define the notion of a possible course of events as well as the notion of momentary open possibilities; see §2.1 below for an overview.

Branching space-times The framework of BT was employed, e.g., for a formal logic of agency (Belnap and Perloff, 1988). Problems in that specific application, as well as considerations of relativistic physics, led to the development of the framework of *branching space-times* (Belnap, 1992), which to date is the most advanced formal framework for describing branching-based possibilities. It allows for the definition of spatiotemporal possible courses of events and truly local possibilities. See §2.2 below for a brief overview.

Branching metaphysics and persons From logic, the notion of branching was introduced into the discussion of modal metaphysics. In fact, Kripke was one of the key proponents of branching-based possibilities. The notion of a rigid designator—a syntactical category of singular expressions, like proper names, that are to refer to the same individual across different possible circumstances—, plays an important role in his *Naming and Necessity* (Kripke, 1980). Kripke’s theory is not limited to proper names for persons; in fact he uses Frege’s example of the morning star and the evening star to make his point about *a posteriori* necessities. However, Kripke does use the example of proper names of persons to make a case against Lewis’s anti-branching, divergence-based counterpart theory for *de re* modality, most famously via his example of U.S. presidential candidate Humphrey, who “couldn’t care less whether someone *else*, no matter how much resembling him, would have been victorious in another possible world” (Kripke, 1980, 45n13). (The point remains contentious; cf., e.g., Hazen (1979) and Mackie (2008).)

Discussions about the metaphysics of personal identity and fission were linked with the idea of branching after Shoemaker’s famous thought-experiment of brain transplants, which was extended to partial brain transplants as possible cases of personal fission; for an overview, see Wiggins (2001). Partially motivated by practical considerations (expectations about the future), Lewis developed his metaphysics of modality based on non-branching (“diverging”) possible worlds as an alternative to a branching theory and published an influential but also much debated argument against branching in his book, *On the plurality of worlds* (Lewis, 1986, 206ff.). The discussion surrounding the metaphysics of branching for individuals, and especially for persons, is intertwined with questions of criteria

for (personal) identity and remains controversial; we will not enter this topic in this paper.

The above examples give strong support to the claim that the uses of the branching metaphor in quantum physics, in logic and in metaphysics, which originated at almost the same time in the late 1950s, are really of a piece: the metaphor, which suggests trunk, directionality, continuity and separation of one into many, works in the same way. The mentioned uses all have a common root in mathematical branching structures and develop such notions in further, domain-specific ways, with the sense of directionality supplied by the temporal order.

2 Formalities of branching

In this section we will be concerned mainly with the exposition of branching time (BT; §2.1), giving only a brief overview and pointers to the literature for branching space-times (BST; §2.2). Both frameworks allow for the definition of *real possibilities*, which are indexically tied to some context rather than abstract. This notion is parallel to the linguists’ category of circumstantial possibility; see, e.g., Kratzer (1991).

2.1 Branching time (BT)

As mentioned above, BT is an application of well-known mathematical structures—backward non-branching partial orders—for the semantics of a temporal-modal language. We will limit ourselves to a simple propositional language that has

- a set *Atoms* of atomic propositions,
- the standard truth-functional connectives, taking conjunction, “&”, and negation, “¬”, as basic and the rest as abbreviations,
- Prior’s tense logical operators for past and future, where we take the strong operators “*H*”, it has always been the case that, and “*G*”, it is always going to be the case that, as basic, and define the weak ones (“*F*”, future, and “*P*”, past) in the standard way:

$$F\phi \Leftrightarrow_{df} \neg G\neg\phi; \quad P\phi \Leftrightarrow_{df} \neg H\neg\phi,$$

- a strong modal operator *Sett*: for “settled truth”, also known as “historical necessity”; the dual weak operator for “historical possibility” is again defined as an abbreviation via

$$Poss:\phi \Leftrightarrow_{df} \neg Sett:\neg\phi.$$

For the semantics, we need the following:

BT-frame: A *BT-frame* is a non-empty partial order $\langle M, \leq \rangle$ (i.e., a nonempty set M with a transitive and antisymmetric relation \leq) such that

- (1) there is no backward branching, i.e., for all $m, m', m'' \in M$ for which $m' < m$ and $m'' < m$, we have either $m' \leq m''$ or $m'' \leq m'$;
- (2) any elements m and m' have a greatest common lower bound in M ;
- (3) M has no maximal or minimal elements, i.e., for any $m \in M$, there are $m', m'' \in M$ s.t. $m' < m$ and $m < m''$.

In clause (1) we have used the strict $<$ order corresponding to \leq , which is defined, as usual, via

$$m < m' \Leftrightarrow_{df} (m \leq m' \ \& \ m \neq m').$$

Of the mentioned conditions, (1) is standard, (2) is sometimes considered optional, but is crucial for our system (guaranteeing connectedness of the structure as well as the right topology in the continuous case), and (3) is just a simplifying assumption saving us a number of extra clauses.

In a BT-frame we can single out the set *Hist* of histories: $h \subseteq M$ is a history iff h is a maximal linear subset, or a maximal chain, in M , i.e., a set that is linearly ordered (for any $m, n \in h$, either $n \leq m$ or $m < n$) and such that no proper superset has that feature. In terms of histories, condition (2) says that all histories have a non-empty intersection, and that that intersection contains a greatest element. By (1), histories are downward closed, i.e., if $m \in h$ and $m' \leq m$, then $m' \in h$.

The basic elements of so-called Ockhamist BT semantics are derived from BT-frames.

BT-model, point of evaluation: The set $M/Hist$ of *points of evaluation* (points at which a formula can be evaluated as to its truth or falsity) consists of all pairs m/h for which $m \in h$, i.e.,¹⁰

$$M/Hist := \{m/h \mid h \in Hist \ \& \ m \in h\}.$$

Thus, generally we need both a moment and a history through that moment to evaluate a formula. A *BT-model* is a BT-frame $\langle M, \leq \rangle$ together with a valuation $V : Atoms \times M/Hist \mapsto \{0, 1\}$.¹¹

Semantics: Given a model $\mathfrak{M} = \langle M, \leq, V \rangle$, the basic clause for atoms is, standardly:

- $\mathfrak{M}, m/h \models p$ iff $V(p, m/h) = 1$.

¹⁰We use the suggestive notation “ m/h ” for the ordered pair of m and h in order to stress the presupposition that $m \in h$.

¹¹Sometimes a valuation is defined as a mapping from $Atoms \times M$ to $\{0, 1\}$; the difference is not important for our discussion here. (Considerations of uniform substitution in a proof system favour the definition given in the main text; see Thomason (1970, 280).)

The clauses for propositional operators are standard. For the tense and modal operators we have the following:

- $\mathfrak{M}, m/h \models H\phi$ iff for all $m' \in h$ for which $m' < m$, we have $\mathfrak{M}, m'/h \models \phi$;
- $\mathfrak{M}, m/h \models G\phi$ iff for all $m' \in h$ for which $m < m'$, we have $\mathfrak{M}, m'/h \models \phi$;
- $\mathfrak{M}, m/h \models Sett:\phi$ iff for all $h' \in Hist$ for which $m \in h'$, we have $\mathfrak{M}, m/h' \models \phi$.

From these clauses we can see that the only operator whose clause really forces the recursive use of a history parameter for evaluation, is the future tense operator G : in the clause for H , the requirement that $m' \in h$ is in fact superfluous, by downward closure of histories, and in the clause for $Sett$;, the initially given history is not used, as all histories through m are quantified over.

This is as is to be expected for a logic for an open, branching future of possibilities: we need to know which of the multiple possible futures we are considering if we want to evaluate a sentence about the future. Crucially, a context of utterance does not supply us with a “history of the context”, which would amount to singling out “the real future”. There is some debate about this feature of the semantics, with some authors proposing a “Thin Red Line” that specifies “the real future” (see, e.g., Øhrstrøm (2009) and Malpass, this issue); there is however a truly indeterministic reply, spelled out by Belnap (2002a). In what follows, we will not pause to consider these semantic matters, but focus instead on a few structural features of BT that should be uncontroversial also for proponents of a Thin Red Line.

The structure of BT suggests a natural definition of the *real possibilities at a moment*, which is based on the notion of *division of histories at a moment*:

- Let m belong to histories h_1 and h_2 , i.e., $m \in h_1 \cap h_2$. h_1 and h_2 are called *undivided at m* (written $h_1 \equiv_m h_2$) iff there is some $m' \in h_1 \cap h_2$ for which $m < m'$. Being undivided at m is an equivalence relation on the set H_m of histories containing m . Reflexivity and symmetry are trivial.¹² For transitivity, let $h_1 \equiv_m h_2$ as witnessed by m' , and $h_2 \equiv_m h_3$ as witnessed by m'' . As $m', m'' \in h_2$, we have $m' \leq m''$ or $m'' \leq m'$ (by linearity of histories); assume the former (the other case is analogous). By backwards linearity we get $m' \in h_3$, so that m' also serves as a witness for $h_1 \equiv_m h_3$.
- We say that h_1 and h_2 *split at m* (written $h_1 \perp_m h_2$) iff m is maximal in $h_1 \cap h_2$.
- The *real possibilities at m* are the members of the partition Π_m of H_m induced by the equivalence relation of undividedness at m , \equiv_m .
- We say that m is a *choice point* iff Π_m has more than one member, i.e., if there are at least two histories splitting at m .

¹²For reflexivity we require that there be no maximal moments; an easy patch is available for the general case.

In this way, BT allows us to talk of real possibilities at a moment: the partition Π_m embodies the real possibilities open at m . These possibilities form an exhaustive set of mutually exclusive alternatives. A full probability theory can be built upon the structure of such possibilities in BT, and even, in BST; see Müller (2005) and Müller (2011b).¹³

The last important feature of BT to be discussed here, which will be important below, is the definability of a global modality. Recall that a guiding idea in modal logic is that possibility is to be truth in some case (or “world”), whereas necessity is truth in *all* cases. This idea was generalized by the now standard Kripke-semantics, which is based on a relation of relative possibility between cases. We can read a BT-frame as a structure of relative accessibility, where the accessibility relation is the (temporal-modal) ordering, \leq .

None of the intensional operators defined above directly expresses “real possibility or necessity *tout court*” (to be symbolized $\diamond\phi$; $\Box\phi$), i.e., the fact that some ϕ holds in *some* or in *all* cases (where the cases are the points of evaluation m/h). And normally we know that we cannot define such a modality by clauses embodying the idea of relative accessibility, like all our semantic clauses above. The basic fact here is that a normal modal logic cannot discern disconnected substructures; modal logic is, as one of the clauses of the Goldblatt-Thomason theorem tells us, invariant w.r.t. disjoint unions of models (Blackburn et al., 2001, 142).

Here however we know, via clause (2) of the above definition of a BT-frame, that we are dealing with a single connected structure, and we have operators moving us back and forth on histories (G and H), as well as an operator that switches histories ($Sett$). This, together with the absence of minima (clause (3) above), allows us the following definition:¹⁴

$$\Box\phi \Leftrightarrow_{df} HSett:GSett:\phi.$$

It is easy to check that this clause really reaches all points of evaluation $m/h \in M/Hist$. The respective global possibility operator is the dual,

$$\diamond\phi \Leftrightarrow_{df} \neg\Box\neg\phi \leftrightarrow PPoss:FPoss:\phi.$$

It is natural to identify this \diamond with the notion of *real possibility tout court*: something is really possible if we can now say that there was a time at which

¹³An alternative, better format for representing such possibilities uses the notion of a *transition* (von Wright, 1963; Belnap, 1999): a transition t is an ordered pair $\langle m, H \rangle$ consisting of an *initial* (a moment $m \in M$), together with an *outcome* (one of the real possibilities at that moment, $H \in \Pi_m$). Employing suggestive notation, we also write $t = m \mapsto H$; TR_m is the set of all transitions with initial m . A transition is *trivial* if there is no alternative transition with the same initial. If m is not a choice point, then $TR_m = \{t\}$ with $t = m \mapsto H_m$, a trivial transition. We can alternatively speak of Π_m or of TR_m as the set of possibilities open at m ; working with sets of transitions simplifies the technical development, e.g., of probability theory. See the references in the main text.

¹⁴In the presence of minima, we can still define the global modality, albeit in a somewhat more involved manner: we can use

$$\Box\phi \Leftrightarrow_{df} Sett:\phi \ \& \ HSett:\phi \ \& \ Sett:G\phi \ \& \ HSett:GSett:\phi.$$

it could have been possible. Such real possibilities are still indexically tied to a context of utterance; they are genuinely thisworldly. As we will see below, other notions of possibility lack this feature.

2.2 Branching space-times (BST)

We remarked above that BST, developed by Belnap (1992), is a natural extension of BT that can be motivated by considerations of agency and/or of relativity theory. Consequently, that framework has been described as “equally proto-scientific and proto-humanistic” (Belnap, 2007a, 19); it allows for applications with respect to both what Sellars (1963) calls “the manifest image” and to what he calls “the scientific image of man”.

The main technical change in moving from BT to BST is that histories in BST are no longer required to be linear; instead, they are taken to be maximal directed sets in a partial ordering (a set X is directed if for any $x, y \in X$ there is some $z \in X$ s.t. $x \leq z$ and $y \leq z$), where that partial ordering is allowed to have backward branching (to be read as spatial separation).

For our purposes here, the most important point is that BST, like BT, allows for the definition of *real possibilities*, which in BST are local in the literal sense of corresponding to possible spatio-temporal courses of events in the causal future of some given event. Again, such possibilities are represented by a partition of a set of histories containing a certain initial event, thereby forming an exhaustive set of pairwise exclusive alternatives. The crucial feature allowing for the definition of the relevant partition is the specific topological structure of branching enforced by the so-called *prior choice principle* of BST, a generalization of condition (2) in the definition of a BT frame above.

In order not to burden this paper with further technical details, we only give references to a number of papers developing the ideas of BST in different directions pointing to specific applications:

- Introductory accounts to BST are given by Belnap (1992) and Belnap (2007a).
- The theory of causation via *causae causantes* as INUS conditions is given in Belnap (2005).
- For an application of BST to the semantics of counterfactuals, see Placek and Müller (2007).
- Objective single case probabilities based on BST are developed in Weiner and Belnap (2006), Müller (2005), and Belnap (2007b). A rather preliminary application in quantum information theory is sketched in Müller (2007); for an analysis of Bell’s theorem, see Placek (2010).
- Minkowskian models for BST, in which each history has the structure of Minkowski space-time, are developed in Müller (2002), Placek and Wroński (2009), and Wroński and Placek (2009).

- The notion of modal correlations, a.k.a. “funny business”, is analyzed in Belnap (2002b), Belnap (2003), and Müller et al. (2008).
- Extensions of the BST framework are considered by Müller (2010), Placek (2011), and Müller (2011a).

3 A landscape of possibilities

We are ready to approach the topic of the paper proper: different notions of possibility, and the place of branching-based possibilities among them. In order to motivate a number of different notions of possibility, we will start by picking up some clues from science, life, and experiment (§3.1). The important issue then is to get an overview of kinds of possibility (§3.2). The question of the role of branching-based real possibility in this taxonomy we postpone to §4 below.

3.1 Motivational

Science seems to tell us, at its current state of development, that the world doesn’t evolve deterministically; there are open future possibilities. This is supported both at the level of phenomena and at the level of theories:

- There are phenomena, stable patterns in empirical data, for which no widely accepted deterministic model is available. Arguably the spreading of an epidemic or the dynamics of populations (both year-to-year and over evolutionary periods of time) are good examples. One may also think of the operational basis of insurance companies. (Those who think that these examples don’t count because they are not fundamental enough, may consider radioactive decay, or the phenomena of quantum correlations; never mind the Bohmians, or superdeterminists.)
- There are well-accepted, highly applicable theories that are not deterministic according to various criteria. Quantum mechanics itself is arguably a good example. Even classical mechanics admits certain forms of indeterminism, e.g., in the form of the famous case of five point masses discussed by Xia (1992). And there is an intricate array of results available for the determinism question in general relativity. See Earman (2007) for details.

It seems therefore that we need a notion of possibility of evolution of the world, or of physical systems. These two don’t necessarily coincide. (In saying this, we are saying that it is possible that they do not coincide. Here we are invoking yet another notion of possibility, which may be identified as conceptual possibility, to be discussed below.)

Moving away from science—closer to home, so to say—we also need a notion, or again various notions, of possibility. Our world is a world of possibilities. Concepts presupposing possibilities are so deeply ingrained in our manifest image, or commonsense view of the world, that it seems hopeless to try and do without.

We hope, we decide, we gamble, praise, blame, and regret missed opportunities. All these practical concepts—and there are many more—presuppose a thick notion of possibility.¹⁵

Well, a number of our practical concepts we can perhaps understand on the basis of an *epistemic*, subjective notion of possibility, where “*X* is possible” just means “*X* is compatible with all I know”. E.g., when we hope for something to turn out a certain way—like I hope that my daughter will lead a happy life when she grows up—this may be taken to mean that it will really necessarily be one way or the other (it is fixed what it will be), but I don’t yet know, and all that the hope involves is a certain preference: that when I finally come to know the facts, I will prefer happy facts over unhappy ones. This may not be a coherent analysis of what hope consists in—hope is a much broader practical concept really, which can, e.g., influence one’s current decisions. But even if we grant an epistemic reading for hope, we surely cannot have such a reading for other practical notions, like regretting a missed opportunity. I could have brought an umbrella, we may think, dripping; I regret not having taken it. It was there on the hatstand for me to take; I really could, and should, have taken it. I could have taken it; it was possible, then, for me to take it. What would an epistemic, subjective reading of that possibility amount to? Before I refrained from taking it, I didn’t know that I wasn’t going to take it? It was compatible with all I knew then that I should take it? That may be so, but something seems to be missing. If before I took off, I didn’t know I wasn’t going to take it, and now I know I didn’t take it, I have simply learned something. That is not something to regret. In fact there is nothing, on the subjective reading, that the specific notion of regret could hook onto.

Clearly we are entering difficult waters here. Some may say, so much the worse for all this folk psychology. We ought not to place undue reliance on the example of regret, since this is part of the notoriously difficult area of action theory. There is no agreement on the analysis of human actions, and in fact the debate seems to be stuck at a dead end. Methodological considerations may have some purchase, but it is hard to say. E.g., a naturalist methodology may urge one to dismiss concepts of folk psychology, while on the other hand, the analytic tradition of conceptual analysis would urge one to start by acknowledging (although in a defeasible way) the natural language concepts that we have.

In order to bridge this methodological gap, it is useful to focus for a while on the notion of a scientific experiment, mainly because for the notion of an experiment, both manifest and scientific image are crucial. On the one hand, experiments more than anything else have helped and are helping to build up the scientific image; a methodological choice to take the sciences seriously in

¹⁵And the same is true, actually, of ordinary theoretical concepts of kinds of things: for something to be a human being, or a king, or a chair, is for indefinitely many *counterfactuals* to be true of it. If it walks like a duck and quacks like a duck, it may still not *be* a duck; we also require that if you took it apart, you’d find a liver and other such things, and that it would cease to exist by being taken apart, even if it never is.

philosophy requires one to take experiment seriously, too.¹⁶ On the other hand, however, experimental work in the lab consists of normal human actions and proceeds, to a large part, in terms of our manifest image. In experiment, we manipulate things, we interfere in natural processes, and all that, ultimately, with our hands and feet. The lab is indeed a tellingly Aristotelian place (see Wiggins (2001, 182)).

There can be no doubt that the notion of a controlled experiment, in which an experimenter “poses a question to nature” by intervening in the course of things to bring about specific experimental initial conditions, has been crucial for the advancement of science. It is a huge difference, well known to practicing scientists, whether one can set up an experiment or one has to rely on observational data from cases whose initial conditions were not controlled.¹⁷ Experiments aren’t always an option, though. It is not always possible to intervene in the course of nature so as to fix certain initial conditions. Here are some quotes, plucked from the internet, that testify to this:

- We can’t do experiments in astronomy, all we can do is observe objects in the universe.
- On the global scale, we can’t do experiments . . .
- We can’t do experiments in large organisms, except in a very few cases.
- We can’t do experiments in humans . . .

These examples can be multiplied easily. As we can see, there are various factors that can make experiments impossible:

- it may be physically impossible to set specific initial conditions: we cannot ionize a single atom such that its total charge is $1/7$ of that of an electron, because charge is quantized;
- bordering on such cases (the line is certainly not strict), it may simply be technically unfeasible to set specific initial conditions: there may be no way really to find out about a snowball’s chances in hell—in a secular environment, let’s say, at the center of the sun—because it is impossible to get it there. Similarly, in cosmology, we can’t do experiments because we can’t move around galaxies, and even if we could, we wouldn’t live long enough (probably not even as a species) to learn anything from that;
- it may be unethical—this plays an important limiting role in experimental physiological or medical research, both on humans and on other animals;

¹⁶Some will deny this, holding that experiment is nothing but observation. This is untenable. For some hard evidence, consider the discussion about the role of the independence of the polarizer settings, a.k.a. “no conspiracy”, in quantum correlation experiments. See, e.g., Peres (1986).

¹⁷We do not need to go as far as Dretske (1994), who proclaims that we only understand nature in so far as we can technically simulate her (“If you can’t make one, you don’t know how it works”). See also note 16.

- it may be unfeasible for economical reasons (we can't pay for it), politically (we don't want to know about this), or the police simply won't let you.

It doesn't seem possible (conceptual possibility again?) to give an epistemic interpretation for any of these notions of possibility—they are all (at least meant to be) objective.¹⁸

3.2 Possibilities: taking stock

Systematizing and taking stock, we have perhaps the following list of possibilities, hopefully doing justice to the image of a “landscape of possibilities” announced in the title. It is certainly not one of the desert landscapes that Quine said he was fond of. (Incidentally, we are not speaking about “modalities” in general, but about “possibilities” specifically. This is not meant to exclude the respective companion notions of impossibility and necessity, but just to narrow down the field to one specific modal dimension. “Modality” is often used very broadly for “all that intensional stuff”, including temporal, deontic, epistemic and other notions. We could do with “alethic modalities”, but that's such a mouthful.) As is customary in logic (see also §2.1 above), we write possibilities as diamonds, with an index to distinguish them. We have:¹⁹

- \diamond_I , *logical* possibility
- \diamond_c , *conceptual* possibility
- \diamond_m , *metaphysical* possibility
- \diamond_n , *natural* (physical) possibility²⁰
- \diamond_t , *technological* possibility

Now we have a few labels; why is that a landscape? What are the interrelations? Let us look at the delineation of these notions of possibility in order to discern some structure. Here are some examples that set off one relation from another:

- \diamond_I (something is green all over and red all over)
- $\neg\diamond_c$ (something is green all over and red all over)
- \diamond_c (Bello is not a dog)

¹⁸Technological possibility is special in that it has a certain subjective element, viz., its historicity: while technological impossibility *now* is intended to be objective, more may become technologically possible (or in fact impossible) later on. Computer typesetting, for example, was technologically impossible a hundred years ago.

¹⁹The list is not meant to be exhaustive. E.g., we omit ethical possibility, even though it also features in the examples of impossibility of experiments mentioned above, because it is not alethic.—One referee suggested to extend the list by a notion of “bureaucratic possibility”, which would have to be given both a time and a place index for determinateness.

²⁰A perhaps helpful idea here is compatibility with the laws of nature (if there are such things). “Physical” may be too narrow, suggesting only compatibility with laws of physics, which is why (following Fine) we prefer “natural”.

- $\neg\Diamond_m$ (Bello is not a dog)
- \Diamond_m (the world is deterministic)
- $\neg\Diamond_n$ (the world is deterministic)
- \Diamond_n (one can travel across the Atlantic in 1h)
- $\neg\Diamond_t$ (one can travel across the Atlantic in 1h)

In these examples it is assumed, of course, that Bello is in fact a dog, that the world is indeterministic as a matter of physical necessity, and that what ballistic intercontinental missiles do doesn't count as traveling.

The examples show that we don't just have different *types* of possibility—there seems to be an *ordering* among them, giving something like layers of possibility, or at least a set-inclusion relation between different sets of possibilities. In fact we have a number of such relations, quite obviously:

- Everything that is possible in any sense, has to be logically possible. Logical impossibility trumps everything else.²¹
- What is technically possible has to be physically possible. We can think of technical possibility as a constraint over and above physical (natural) possibility: for you to be able to build something, it has to be physically possible *and* feasible.

That's fair enough, but what about the real dependence among these notions? There is some debate about this. Do we need to acknowledge several irreducible kinds of possibility, or will a single kind do? Can we defend a monism of possibilities?

The logical empiricists thought that if you had to acknowledge any kind of possibility at all, it had better be logical possibility, and the hope (e.g., Carnap's hope) was that one could analyze that notion in purely syntactic terms. But that

²¹Identity is a logical notion, and it is logically impossible for there to be a , b and c so that $a = b$, $a = c$ but $b \neq c$. As a direct corollary we can dismiss most if not all stories about personal fission as not being based on any possibilities and thus as useless for philosophical argumentation. (Note that cases of "split personality", whether induced via severing the *corpus callosum* or as a result of mental disorder, cannot count as personal fission either: the relation of a brain a giving rise to a consciousness b and a consciousness c is certainly not identity, again on pain of contradiction.) See Wiggins (2001, 72) for an elaboration of this point in the amoeba case, where fission as a means of persistence seems to have at least some *prima facie* plausibility.

There is a much deeper question behind this, which we cannot hope to answer in this paper: How come that we seem to *understand* stories that violate logical possibility? What do we see, e.g., when we see a picture with impossible shadows, or two Captain Kirks in the Star Trek episode of the malfunctioning teleporter? We broadly agree with van Inwagen (1993, 229) that "one may not use examples from fantasy in conceptual investigations [...] [because] the author of a fantasy has the power to confer 'truth in the story' on known conceptual falsehoods". It seems very hard, however, to diagnose what is really going on when we are dealing with such stories. Maybe the pretence theory of fiction (Walton, 1990), which has recently been used in an interesting account of scientific modeling (Frigg, 2010), can be of use here. The question undoubtedly merits further study.

program has failed—it is widely recognized in philosophy today that stronger notions of possibility are (also) needed. Present-day monism about possibilities needs a different basis. Looking at the literature it seems that *metaphysical possibility* is the best candidate for a basic modality. Or is it? Natural (physical) possibility may also be a contender.

There is a convincing argument by Kit Fine showing that we will not be able to reduce the one to the other either way. Fine (2005) starts by pointing out that when we have two notions of possibility ordered by inclusion, as in the case under discussion, there are two strategies available, depending on where we choose to start.²² On the one hand, one can try to define the broader notion, metaphysical possibility, as a *widening* of the narrower notion, natural possibility. According to this recipe, ϕ is metaphysically possible exactly if it is either (already) naturally possible, or it is (actually) a metaphysical truth. On the other hand, one can try to define the narrower notion, natural possibility, as a *restriction* of the broader notion of metaphysical possibility. Thus, assume that the non-modal natural truths are given as a proposition P . (Never mind the infinite case; we're en route to a *reductio*. Let's *give* the opponent finitude.) Then ϕ is naturally possible if it is metaphysically possible that $\phi \wedge P$; i.e., there has to be a metaphysically possible scenario in which both ϕ and all non-modal natural facts hold true. There appear to be no other plausible means of reduction.

Will this work? Consider widening first. The given definition will fail immediately if there are any cases of genuine metaphysical possibilities that are neither truths of metaphysics nor natural possibilities. And that seems to be so: There are metaphysical mere possibilities that aren't also physical possibilities. Fine's actual argument is rather fine-spun, accommodating post-Kripkean worries about alien properties and about the existence of universals in worlds in which they are not instantiated (Fine, 2005, 240f.). But the following will do: it is certainly both a metaphysical possibility that the world is deterministic, and that it is indeterministic. But supposedly the world is one way or the other, and its being so (of course, indeterministic) will be a natural necessity. So here we have a metaphysical mere possibility (a metaphysical proposition that is in fact false, but could be true)—that the world is deterministic—that is not a physical possibility. See Fine (2005, 241).

In the other direction, restriction, we should grant that every physical possibility is also a metaphysical possibility. But what about P , the set of natural facts that need to be available in order to tell the (fewer) natural possibilities apart from the (larger group of) metaphysical possibilities? Recall that ϕ is to be naturally possible just in case $\phi \wedge P$ is metaphysically possible. The worry is that any P will be too weak to narrow down the range of metaphysical pos-

²²Fine gives his argument in terms of the corresponding necessities; we are transposing it to the case of possibilities. We therefore have to alter the terminology. Fine says that the two available strategies for reducing natural and metaphysical necessity to one type of necessity are *restriction*, reducing the broader notion of natural necessity to the narrower one of metaphysical necessity, and *relativization*, reducing the narrower notion of metaphysical necessity to the broader one of natural necessity.

sibilities. Take any set P of occurrent natural facts; how is that supposed to exclude (render *metaphysically* impossible) a *de facto* naturally impossible ϕ ? How could this work without admitting a basic notion of physical possibility through the back door? A more general worry is that it will not be possible anyway to show that modal facts—say, laws of nature—supervene on non-modal ones—say, the Humean mosaic that Lewis bases his Humean supervenience program on. Consider two metaphysically possible worlds that differ with respect to their laws. Now take out all the bodies—that should be a metaphysical possibility as well. We’re left with two empty worlds, surely indistinguishable in terms of Humean properties; but by assumption, these worlds harbour different modal facts. So, facts about natural modality do not supervene on the non-modal natural facts: there is irreducible natural modality. (There is no need to worry about the simplistic idea of empty worlds playing a decisive role in this argument; a more sophisticated version can be made to work as well. See Fine (2005, 244f.).—This is not meant to suggest that everyone takes the debate to be settled.)

So it seems that we are stuck with at least two basic kinds of possibility. (Actually Fine argues for three basic kinds, a further one being the counterpart of normative necessity. We will not discuss that notion here; see also note 19.) Maybe we can reduce the rest—logical possibility, technical possibility and so forth—to those two. (Fine suggests that that may be doable.) But there seems to be no hope for a monism of possibilities. Our landscape really has some distinguishing features.

The question that arises now is: what is the role of branching-based real possibilities vis-à-vis the landscape of possibilities sketched so far?

4 Branching-based possibilities as basic

The upshot of our discussion in §3 was that possibilities are important and intricate, and that it makes sense to speak of a landscape of possibilities because there are several different kinds. This leads to our crucial question: what about branching? How do the branching-based real possibilities that were introduced formally in §2, relate to this discussion; where is a place for them in our landscape of possibilities?

In this section we will try to articulate this question in more detail (§4.1) before we argue, first (§4.2), that we cannot even understand the mentioned notions of possibilities unless we acknowledge the branching-based variety. That is an *epistemic* point, which seems to be reasonably secure. Secondly, and more speculatively, we will argue that branching-based real possibilities are *metaphysically* basic as well (§4.3). So *pace* Fine, we might be able to defend a specific type of monism of possibilities after all. This would however not turn our landscape into a desert—rather, it would show that the richness of the landscape of possibilities is, in an important sense, the same as the richness of the world we live in.

4.1 The question of the place of branching-based possibilities

Recall that we singled out branching possibilities formally: they should be such that they can be fruitfully modeled in terms of branching histories, either along the lines of BT’s idea of linear histories, or in a more sophisticated, BST-type way. The accompanying metaphysical image behind branching is the notion of an open future of possibilities in contrast to the settledness of the past: once open possibilities become unavailable, or drop off, as time progresses. Does this notion of possibility correspond to any of the notions discussed in §3.2 above? Or is there a radical difference?

We often invoke real possibilities with their specific temporal features, especially in practical contexts (see the example of regret above). They occupy an important place for concrete issues in our lives: choices, decisions, and expectations. Accordingly, we may be tempted to draw a distinction according to which real possibilities are appropriate for all *concrete* situations—ones that we can identify *indexically*, or point to, so to speak—, whereas the other notions of possibility are always *abstract*, not tied to concrete situations, and pertain to situations only *as described* qualitatively one way or the other, without any link to temporality. On this view, there would be a *formal* difference between branching-based possibilities and our abstract kinds of (e.g.) logical, physical, or metaphysical possibility.²³

Indeed, branching possibilities are always tied to a context indexically, and the branching possibility operator interacts with other indexicals and their respective modalities in a specific way. Reconsider the regret case from above. I could have taken my umbrella then: it was the case in a concrete past situation that I could take it. Now that is no longer possible; I left the house without the umbrella and got soaked. We can write this up in the following way (using Prior’s P for the past tense):

$$P\Diamond\phi \wedge \neg\Diamond\phi,$$

it was the case that ϕ was possible but (now) it isn’t possible (any more). We know that this is the way our lives go. There are all these possibilities, and we lose more and more of them as we grow older.²⁴ This feature is absent from all abstract types of possibility: if a situation is given abstractly, via an atemporal description, then if it can be assessed with respect to its objective possibility or impossibility at all, that assessment cannot change over time.²⁵ Accordingly, the

²³There may be a further *metaphysical* difference in that branching-based possibilities are grounded in actually existing things—substances—and their modal properties: dispositions, powers, potentialities. We cannot explore this line of thought any further here, and we will therefore continue to take possibilities to be basic for what follows.

²⁴If there are infinitely many possibilities, their cardinality may remain constant—but still, the set of possibilities diminishes over time.

²⁵The case of technological possibility, on whose temporality we have commented above (see note 18), is subtle in this respect. There can be a case in which something that was technologically possible earlier on, becomes technologically impossible later, e.g., due to the depletion of natural resources. (I am indebted to Gerhard Schurz for this point.) Let us

formal demarcation is sharp and determinate. But phenomenologically, the issue is not so clear. We also use the more abstract notions of physical, metaphysical, and other possibilities discussed above in talking about very concrete, indexically identifiable situations. Indeed, referring to the physical impossibility of some concrete thing is often used as a valid excuse or explanation. Here are a few random quotes from the internet that support the point:

- [I]t was physically impossible for me to travel from Batangas City to Makati in time for the event.
- It was physically impossible for me to be at every different field trip and event.
- [I]t was physically impossible for me to make the deadline (since you also needed to mail the UPC).
- I wanted to attend some talks hosted by universities in the US, but it was physically impossible for me to participate.
- It was technically impossible for me to include the ppt file to my website.

One also finds more contentious uses, like “It is physically impossible for me to listen to Rush”, in which the modal claim may function merely as emphasis. But it seems clear that at least the notion of physical possibility is sometimes used when referring to concrete situations, not just to abstract scenarios.

Here is an idea on how this might be explained. We have seen above that the notion of real possibility *tout court*, \diamond , definable via the primitive modalities P , F and $Poss$: of branching time, captures our whole tree of moments and histories. Saying that something is really possible means that there is a witness—a real possibility—that is linked indexically to the moment of utterance. If we use one of the other notions of possibility, then there are two options:

- (1) It may be that what is said to be X -possible—physically possible, for example—is in fact really possible. This means that talking about a physical possibility in such a case may just be a gloss on the manner of insight into the possibility. The role of the available witness may still be left open. Thus, it may be that we do not have to trace histories back and forth, but just consider some physical facts: we may not even care that that which we call physically possible, is additionally *really* possible.
- (2) What is said to be, e.g., physically possible, may be *merely* physically, but not really possible. Thus, there may be no real possibility, from the beginning of our universe to the end of time, that two spheres of solid gold

hope that it will not be the case that in 100 years’ time, our descendants can say that it was technologically possible to stop global warming in 2011, but it is no longer possible in 2111. Still, the case about the abstractness applies if we read the time-index as belonging to the possibility operator itself, so that in the example, we are in fact dealing with two different operators, technological-possibility-in-2011 and technological-possibility-in-2111. Both of them then apply to abstractly specified scenarios, without specific links to temporality.

exactly the size of the moon collide head-on at a relative speed of 517 m/s; but surely such a thing is physically possible, and we can even calculate fairly accurately what might happen as a result. In this case, we have to transcend our homely, indexically connected, branching notion of possibility and jump, so to speak, to a merely hypothetical universe where such a thing happens.

In case (1), the abstract notions of X -possibility play the role of classifying various groups of real possibilities, so they do not actually point to a metaphysically distinct notion of possibility. In case (2), however, we do need a notion of mere X -possibility (an *unreal possibility*, as we might call it), where the merely possible scenarios are not connected to us indexically, but are given via atemporal descriptions.

The notion of mere X -possibilities surely does useful work for physical, metaphysical, logical, and other considerations. Most exercises in a physics textbook, for example, are about merely physically possible scenarios, and yet solving them is the way for us to learn the trade, even if in the end our aim may be to apply it in thisworldly situations. On the other hand, one should not underestimate the range of real possibilities. There is a substantial amount of real possibilities, and once we recognise them, we are free to classify them more or less in any way we please, e.g., by abstracting from specific features. Given a certain practical mindset, it may be useful to classify certain real possibilities as dangerous or safe; in a more theoretical mood, we may classify some as physical (we understand why they are possible since we know the physics that generates them) and others as conceptual (maybe because we can picture them clearly, without having any idea about their physical basis). In this way, we can view the abstract notions of possibility discussed above as abstractions or generalizations arising out of underlying branching-based real possibilities.

This has an important consequence: If real possibilities are what gives rise to our landscape of abstract possibilities, then what is really possible has to be possible in every abstract sense of the word. It can't be that something is really possible while it is, e.g., conceptually or physically impossible.²⁶ Formally, we thus have

$$\diamond\phi \rightarrow \diamond_X\phi,$$

where \diamond stands for real possibility *tout court*, and \diamond_X is any one of the abstract notions of possibility mentioned in §3.2 above.²⁷ The other direction needn't hold—in fact, it provides us with a test for checking whether we are dealing

²⁶The usual *caveats* with respect to the temporality of technical possibility apply; see note 25 above.—To comment briefly on a further issue: Of course we can *say* that something that is really possible, is impossible in some abstract sense. Bumblebees fly, that one over there is flying, so it's really possible for that one to fly; actuality implies real possibility of course. But we hear it said that it is technically impossible for bumblebees to fly. This can probably only mean that there isn't yet a good explanation for the fact that they fly, and that we cannot (yet?) build a machine that would do the same. (And actually even that is a myth: what bumblebees can't do, really, is *glide*. See Zetie (1996) for an illuminating account.)

²⁷Again, the temporality of technical possibility (see note 25 above) requires a qualification: the appropriate X has to be technological possibility *now*; not, say, 200 years ago.

with case (1) or (2) from the above list. Formally, case (1), X -possibility that is witnessed by real possibility, amounts to

$$\Diamond_X\phi \ \& \ \Diamond\phi,$$

while case (2), mere X -possibility, is signalled by

$$\Diamond_X\phi \ \& \ \neg\Diamond\phi,$$

i.e.,

$$\Diamond_X\phi \ \& \ \neg Poss:FPoss:\phi.$$

4.2 The fundamental epistemic role of branching-based possibilities

How can we learn about possibilities? Empiricist worries surround the notion of modal knowledge, but denying that we have such knowledge is really just a variant of skepticism. Minimally, we know that everything that really happened, was really possible before it happened; we know that actuality implies possibility:

$$\phi \rightarrow (P\Diamond F\phi \ \& \ \Diamond_X\phi).$$

That, however, does not give us an interesting notion of modal knowledge yet; a skeptic could claim that the possible simply coincides with the actual, so that we have modal knowledge only to the extent to which we have non-modal knowledge. If one wants to avoid skepticism, one needs to argue that we have knowledge of *unrealized* possibilities as well. But that, too, seems to be beyond reasonable doubt. Any real case of regret testifies to an unrealized real possibility, and as noted above (the safer example), even our theoretical concepts of kinds of objects presuppose modality (see note 15). There are countless known counterfactual truths about the ordinary objects surrounding us. Most of these are instances of knowledge of real possibility. But then, we also know a vast number of possibilities and impossibilities that are not grounded in real possibility, but abstract. For example, we can acquire a lot of knowledge of logical possibilities by simple formal considerations—so many in fact that, as mentioned above, for a time in the history of logical empiricism it seemed that all useful knowledge of possibility could stem from this formal source.

We can take it as a given that we have various sorts of modal knowledge. The question we wish to ponder here is whether knowledge of one sort of possibilities can count as basic, or whether there are independent epistemic stories to be told about our knowledge of real vs. unreal possibilities. The claim we wish to defend is that knowledge of real possibilities is the basic phenomenon, on which knowledge of unreal possibilities depends.

Research on causal knowledge has made it pretty clear that as human beings we have the capability to build up modal knowledge from a very early age (Leslie and Keeble, 1987; Gopnik and Schulz, 2007; Dullstein, 2008). Both observation and active manipulation of the world around us play an important role in this

process. It seems reasonable to assume that if we could not actively intervene in the course of things (like Dummett's imagined society of intelligent trees (Dummett, 1964)), our knowledge of real possibilities would be severely limited.

The fact that we acquire modal knowledge of objects surrounding us, both via observation and by active manipulation, underwrites our claim. We build up knowledge of concrete possibilities and impossibilities. We often know whether we can move a certain object before us, we may have tried before successfully or unsuccessfully. I know that the cup before me might fall and spill coffee over the keyboard; if I were to push it hard enough from *this* direction, it would fall. (I have spilled coffee before.)

The importance of manipulation provides an illuminating view on the epistemological function of scientific experiments, and on the role of experiment in establishing abstract modal facts (including, e.g., facts about the merely physically possible scenarios of our physics textbooks). In an experiment, we exploit real possibilities in the lab, setting initial conditions and actively manipulating things in order to pose a specific question to nature. We interact with the world in a concrete situation in the lab; from our practice we know of real possibilities that would mean that the experiment was a failure (e.g., insufficient shielding from outside influences, or, in a biological case quoted by Bogen and Woodward (1988), the scientist's boss's heavy steps on the staircase). Suppose that a result is read off after the experiment. That reading was, before the experiment, a real possibility. Depending very much on the local facts, that real possibility may be interpreted as a *physical* possibility for an abstractly described (repeatable) experimental set-up. The relevant step of inference has much in common with, and may in fact be part of, the establishment of phenomena from data (Bogen and Woodward, 1988). At the relevant level of abstraction, concrete real possibilities that are too "local", indicating various failures of the experiment, will be abstracted from. In such a way the experimenter, through her local modal knowledge, escapes the experimenter's regress. The abstract modal facts gained through experimentation can then be used to derive general modal facts describing the behaviour of classes of systems, up to the formulation of laws of nature.

As agents, we have a vast amount of knowledge of real possibilities before we are ever in a position to even grasp the abstract concepts of unreal possibilities discussed above. They are clearly epistemically secondary; modelled, perhaps, on some formal properties of real possibilities, or on general rules of games of make-believe. Once we have them, however, the table may be turned: If I know that the situation before me is of a type for which it is X -possible that ϕ , and there are no overriding factors (the infamous so-called *ceteris paribus* clause), then I may well be in a position to argue that it is, in the present situation, *really* possible that ϕ . What we have said here is not meant to undermine the relevance of such inferences, but just to express a doubt as to their basic epistemological status, and to indicate a way for understanding the *ceteris paribus* clause.

We have to leave a more detailed discussion of the epistemology of real vs. unreal possibilities for another occasion. In concluding, we want to comment at least very briefly on the metaphysical side of the distinction.

4.3 A fundamental metaphysical role of branching-based possibilities?

If we accept that we have knowledge of real possibilities—and above we have argued that a denial amounts to a form of skepticism—, then clearly *there are* these real, branching-based possibilities for us to have knowledge about. Our world, really, is a world of possibilities, of things and living beings with their individual modal properties. In a branching world, there is room for substantial thisworldly modality—and such a world is a world even according to the definition of Lewis (1986, 208), who demands unification by “suitable external relations”, which are given in the form of the temporal-modal ordering, \leq .

This line of argument establishes a basic metaphysical role of real possibilities, and our comments about their epistemology should at least make plausible that we can have knowledge about them. What about the unreal possibilities that figured so prominently in our first sketch of a landscape of possibilities, and which pervade the philosophical discussion of modality?

A simple answer, which fits rather well into our overall picture, is that unreal possibilities are just that: unreal. So there are, in an important sense, no unreal possibilities. They are not there for us to use, to get to know by trying things out, or to do anything with—except for the classificatory use, grouping together different sets of real possibilities, that we acknowledged above, according to which they are mere “creatures of the understanding” (to use a Lockean phrase).

But what of our landscape, then? What about the role of abstract, unreal possibilities in physics, in logic, in philosophical argumentation and elsewhere? We certainly do not want to deny, or to downplay, the usefulness of our landscape of abstract possibilities.

The image of possible worlds may be genuinely helpful at this juncture—but we will use it in a substantially transformed way. In Lewis’s usage, possible worlds are there to do all the modal work in a unified fashion, providing the basis for the most minute *de re* modal truths (“I could have one hair less than I actually have”) to the wildest counterfactuals (“If there was time travel, humankind could win the robot wars”). Lewis achieves this unification by investing heavily into a context-dependent notion of similarity between possible worlds. This approach, which parallels the linguistic theory of Kratzer (1991), may be good for the analysis of communication, in the context of which it was developed. What we propose, however, is that in the application to metaphysics, the unified machinery misconstrues a crucial ontological distinction—that between real possibilities, which are genuinely thisworldly, and unreal possibilities, which are not—as a mere matter of context dependence.

We offer a different picture, and our claim is that our package is, on the whole, more attractive. It is less revisionary, based on a clear formal semantics, and keeps in touch with scientific practice. On our picture, the ontological distinction between branching-based real possibilities and merely abstract possibilities coincides with the distinction between thisworldly and otherworldly possibilities. Our world, the world in which we live, the world we interact

with, is rich enough to harbour a whole landscape of real possibilities, based on the things around us. In contrast, we can think of unreal possibilities in the manner of *separate* possible worlds, not connected to us via “suitable external relations”. Physics textbooks may instruct us to imagine barren worlds of two classical point particles, and to solve their equations of motion. The analysis of an argument may point us towards abstract situations in which all premisses are true, and ask us to determine whether the conclusion is true “there”, too. We certainly need to take these moves seriously—but there is no need for us to put serious metaphysical weight on them: this may be games of make-believe.²⁸ *Pace* Lewis’s modal realism, a good case can be made for a radical ontological distinction between our world and the worlds of unreal possibilities. Remember, our world itself is replete with possibilities. These may be all that is needed for our practical concerns, including our experimental work in the lab.

I could really write on, but the journal won’t let me.

Acknowledgements

I would like to thank audiences in Bristol and Utrecht for many helpful discussions, and two referees for their comments and criticism. Support by VIDI grant 276-20-013 of the Netherlands Organization for Scientific Research is gratefully acknowledged. Special thanks to Sebastian Lutz and Jesse Mulder for their detailed feedback on a previous draft.

References

- AMS: 2010, ‘American Mathematical Society Mathematics Subject Classification 2010’. URL = <http://www.ams.org/mathscinet/msc/>.
- Belnap, N.: 1992, ‘Branching space-time’. *Synthese* **92**(3), 385–434. See also the postprint 2003, available on [philsci-archive](http://philsci-archive.net/).
- Belnap, N.: 1999, ‘Concrete Transitions’. In: G. Meggle (ed.): *Actions, Norms, Values: Discussions with Georg Henrik von Wright*. Berlin: de Gruyter, pp. 227–236.
- 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, pp. 37–58.
- Belnap, N.: 2002b, ‘EPR-like “funny business” in the theory of branching spacetimes’. In: T. Placek and J. Butterfield (eds.): *Non-locality and Modality*. Dordrecht: Kluwer, pp. 293–315.

²⁸See note 21 above for a link between make-belief and scientific modeling.

- Belnap, N.: 2003, 'No-common-cause EPR-like funny business in branching space-times'. *Philosophical Studies* **114**, 199–221.
- Belnap, N.: 2005, 'A theory of causation: *Causae causantes* (originating causes) as inus conditions in branching space-times'. *British Journal for the Philosophy of Science* **56**, 221–253.
- Belnap, N.: 2007a, 'From Newtonian determinism to branching space-time indeterminism'. In: T. Müller and A. Newen (eds.): *Logik, Begriffe, Prinzipien des Handelns*. Paderborn: Mentis, pp. 13–31.
- Belnap, N.: 2007b, 'Propensities and probabilities'. *Studies in History and Philosophy of Modern Physics* **38**(3), 593–625.
- Belnap, N. and M. Perloff: 1988, 'Seeing to it that: a canonical form for agents'. *Theoria (Lund)* **54**(3), 175–199.
- Blackburn, P., M. De Rijke, and Y. Venema: 2001, *Modal logic*. Cambridge: Cambridge University Press.
- Bogen, J. and J. Woodward: 1988, 'Saving the phenomena'. *The Philosophical Review* **97**(3), 303–352.
- Byrne, P.: 2010, 'Everett and Wheeler: the untold story'. In Saunders et al. (2010), pp. 521–541.
- Church, A.: 1956, *Introduction to Mathematical Logic, Vol. I*. Princeton, NJ: Princeton University Press.
- Cooper, L. N.: 1976, 'How Possible Becomes Actual in the Quantum Theory'. *Proceedings of the American Philosophical Society* **120**(1), 37–45.
- Cooper, L. N. and D. Van Vechten: 1969, 'On the interpretation of measurement within the quantum theory'. *American Journal of Physics* **37**, 1212–1220.
- DeWitt, B. S.: 1970, 'Quantum mechanics and reality'. *Physics Today* **23**(9), 30–35.
- Dretske, F.: 1994, 'If you can't make one, you don't know how it works'. *Midwest Studies in Philosophy* **19**, 468–482.
- Dullstein, M.: 2008, 'Verursachung und kausale Relevanz. Eine Analyse singulärer Kausalaussagen'. Ph.D. thesis, University of Heidelberg.
- Dummett, M.: 1964, 'Bringing about the past'. *Philosophical Review* **73**(3), 338–359.
- Earman, J.: 2007, 'Aspects of Determinism in Modern Physics'. In: J. Butterfield and J. Earman (eds.): *Handbook of the Philosophy of Physics*. Amsterdam: Elsevier, pp. 1369–1434.

- Everett: 1957, ‘Relative state’ formulation of quantum mechanics’. *Rev. mod. phys.* **29**(3), 454–462.
- Feathers, N.: 1947, ‘Radioactive branching’. *Reports on Progress in Physics* **11**, 19–31.
- Fine, K.: 2005, *Modality and Tense*. Oxford: Oxford University Press.
- Frigg, R.: 2010, ‘Models and fiction’. *Synthese* **172**, 251–268.
- Gopnik, A. and L. Schulz (eds.): 2007, *Causal Learning. Psychology, Philosophy, and Computation*. Oxford: Oxford University Press.
- Hazen, A.: 1979, ‘Counterpart-theoretic semantics for modal logic’. *Journal of Philosophy* **76**(6), 319–338.
- Hopf, E.: 1942, ‘Abzweigung einer periodischen Lösung von einer stationären Lösung eines Differentialsystems’. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig. Mathematisch-Physikalische Klasse* **94**, 3–22.
- Kolmogorov, A. N. and N. A. Dmitriev: 1947, ‘Branching stochastic processes’. *Comptes Rendus (Doklady) de l’Académie des Sciences de l’URSS* **56**(1), 5–8. See also *Math Rev* 9,1 (1948) 46.
- Kratzer, A.: 1991, ‘Modality’. In: A. von Stechow and D. Wunderlich (eds.): *Semantics: An international handbook of contemporary research*. Berlin: de Gruyter, pp. 639–650.
- Kripke, S.: 1980, *Naming and Necessity*. Oxford: Blackwell. Originally published in 1972.
- Leslie, A. M. and S. Keeble: 1987, ‘Do six-month-old infants perceive causality?’. *Cognition* **25**, 265–288.
- Lewis, D. K.: 1986, *On the plurality of worlds*. Oxford: Blackwell.
- Lindenbaum, A. and A. Tarski: 1936, ‘Über die Beschränktheit der Ausdrucksmittel deduktiver Theorien’. *Ergebnisse eines mathematischen Kolloquiums* **7**, 15–22. English translation in Tarski, *Logic, Semantics, Metamathematics*, Hackett 1983, Ch. XIII.
- Mackie, P.: 2008, ‘Transworld Identity’. In: E. N. Zalta (ed.): *The Stanford Encyclopedia of Philosophy*. Winter 2008 edition.
- Müller, T.: 2002, ‘Branching space-time, modal logic and the counterfactual conditional’. In: T. Placek and J. Butterfield (eds.): *Non-locality and Modality*. Dordrecht: Kluwer, pp. 273–291.
- Müller, T.: 2005, ‘Probability theory and causation: a Branching Space-Times analysis’. *British Journal for the Philosophy of Science* **56**(3), 487–520.

- Müller, T.: 2007, 'A branching space-times view on quantum error correction'. *Studies in History and Philosophy of Modern Physics* **38**, 635–652.
- Müller, T.: 2010, 'Towards a theory of limited indeterminism in branching space-times'. *Journal of Philosophical Logic* **39**, 395–423.
- Müller, T.: 2011a, 'Branching space-times and the Hausdorff property'. Forthcoming.
- Müller, T.: 2011b, 'Probabilities in branching structures'. In: F. Stadler (ed.): *The Philosophy of Science in a European Perspective, Vol. 2*. Dordrecht: Springer. In print.
- Müller, T., N. Belnap, and K. Kishida: 2008, 'Funny business in branching space-times: infinite modal correlations'. *Synthese* **164**, 141–159.
- Øhrstrøm, P.: 2009, 'In defence of the Thin Red Line: a case for Ockhamism'. *Humana.mente* **8**, 17–32.
- Peres, A.: 1986, 'Existence of "free will" as a problem of physics'. *Foundations of Physics* **16**(6), 573–584.
- Placek, T.: 2010, 'Bell-type correlations in branching space-times'. In: T. Czarnecki, K. Kijania-Placek, O. Poller, and J. Woleński (eds.): *The Analytic Way. Proceedings of the European Congress on Analytic Philosophy, Kraków, August 2008*. London: College Publications.
- Placek, T.: 2011, 'Possibilities without possible worlds/histories'. *Journal of Philosophical Logic* p. forthcoming.
- Placek, T. and T. Müller: 2007, 'Counterfactuals and historical possibility'. *Synthese* **154**, 173–197.
- Placek, T. and L. Wroński: 2009, 'On infinite EPR-like correlations'. *Synthese* **167**(1), 1–32.
- Prior, A. N.: 1957, *Time and Modality*. Oxford: Oxford University Press.
- Prior, A. N.: 1962, 'Possible worlds'. *The Philosophical Quarterly* **12**(46), 36–43.
- Russell, B.: 1908, 'Mathematical logic as based on the theory of types'. *American journal of mathematics* **30**(3), 222–262.
- Saunders, S.: 2010, 'Chance in the Everett interpretation'. In Saunders et al. (2010), pp. 181–205.
- Saunders, S., J. Barrett, A. Kent, and D. Wallace (eds.): 2010, *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford: Oxford University Press.
- Sellars, W.: 1963, *Science, perception and reality*, Chapt. Philosophy and the scientific image of man, pp. 1–40. London: Routledge & Kegan Paul.

- Thomason, R. H.: 1970, 'Indeterminist time and truth-value gaps'. *Theoria* **36**, 264–281.
- Thomason, R. H.: 1984, 'Combinations of tense and modality'. In: D. Gabbay and G. Guentner (eds.): *Handbook of philosophical logic, vol. II: extensions of classical logic*, Vol. 165 of *Synthese Library, Studies in epistemology*. Dordrecht: D. Reidel Publishing Company, pp. 135–165.
- van Inwagen, P.: 1993, 'Naive mereology, admissible valuations, and other matters'. *Noûs* **27**(2), 229–234.
- von Wright, G. H.: 1963, *Norm and Action. A Logical Inquiry*. London: Routledge.
- Walton, K. L.: 1990, *Mimesis as make-believe: On the foundations of the representational arts*. Cambridge MA: Harvard University Press.
- Waring, R. S.: 1882, 'Method of dividing and branching electric cables'. Patent 268 324, Pittsburgh, 28 November 1882.
- Weiner, M. and N. Belnap: 2006, 'How causal probabilities might fit into our objectively indeterministic world'. *Synthese* **149**, 1–36.
- Whitehead, A. N. and B. Russell: 1910, *Principia Mathematica*. Cambridge: Cambridge University Press.
- Wiggins, D.: 2001, *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.
- Wroński, L. and T. Placek: 2009, 'On Minkowskian Branching Structures'. *Studies in History and Philosophy of Modern Physics* **40**, 251–258.
- Xia, Z.: 1992, 'The existence of noncollision singularities in Newtonian systems'. *Annals of Mathematics* **135**(3), 411–468.
- Zetie, K.: 1996, 'The strange case of the bumble-bee that flew'. *Physics World* **9**(10), 72.