

Introduction: The Many Branches of Belnap’s Logic

Thomas Müller

Abstract In this introduction to the *Outstanding contributions to logic* volume devoted to Nuel Belnap’s work on indeterminism and free action, we provide a brief overview of some of the formal frameworks and methods involved in Belnap’s work on these topics: theories of branching histories, specifically “branching time” and “branching space-times”, the *stit* (“seeing to it that”) logic of agency, and case-intensional first order logic. We also draw some connections to the contributions included in this volume. Abstracts of these contributions are included as an appendix.

Nuel Belnap’s work in logic and in philosophy spans a period of over half a century. During this time, he has followed a number of different research lines, most of them over a period of many years or decades, and often in close collaboration with other researchers:¹ relevance logic, a long term project starting from a collaboration with Alan Anderson dating back to the late 1950s and continued with Robert Meyer and Michael Dunn into the 1990s; the logic of questions, developed with Thomas Steel in the 1960s and 1970s; display logic in the 1980s and 1990s; the revision theory of truth, with Anil Gupta, in the 1990s; and a long-term, continuing interest in indeterminism and free action. This book is devoted to Belnap’s work on the latter two topics. In this introduction, we provide a brief overview of some of the formal frameworks and methods involved in that work, and we draw some connections to the contributions included in this volume. Abstracts of these contributions are presented in Appendix A.

¹ The biographical interview with Nuel Belnap provides some additional information on these research lines and on some of the collaborations.

T. Müller (✉)
Department of Philosophy, Utrecht University, Janskerkhof 13a,
3512 BL Utrecht, The Netherlands
e-mail: Thomas.Mueller@phil.uu.nl

1 About this Book

This book contains essays devoted to Nuel Belnap’s work on indeterminism and free action. Philosophically, these topics can seem far apart; they belong to different sub-disciplines, viz., metaphysics and action theory. This separation is visible in philosophical logic as well: The philosophical topic of indeterminism, or of the open future, has triggered research in modal, temporal and many-valued logic; the philosophical topic of agency, on the other hand, has led to research on logics of causation and action. In Belnap’s logical work, however, indeterminism and free agency are intimately linked, testifying to their philosophical interconnectedness.

Starting in the 1980s, Belnap developed theories of indeterminism in terms of branching histories, most notably “branching time” and his own “branching space-times”. At the same time, he pursued the project of a logic of (multi-)agency, under the heading of *stit*, or “seeing to it that”. These two developments are linked both formally and genetically. The *stit* logic of agency is built upon a theory of branching histories—initially, on the Prior-Thomason theory of so-called branching time. The spatio-temporal refinement of that theory, branching space-times, in turn incorporates insights from the formal modeling of agency. Both research lines arise in one unified context and exert strong influences on each other.²

This volume appears in the series *Outstanding contributions to logic* and celebrates Nuel Belnap’s work on the topics of indeterminism and free action. It consists of a selection of original research papers developing philosophical and technical issues connected with Belnap’s work in these areas. Some contributions take the form of critical discussions of his published work, some develop points made in his publications in new directions, and some provide additional insights on the topics of indeterminism and free action. Nearly all of the papers were presented at an international workshop with Nuel Belnap in Utrecht, The Netherlands, in June 2012, which provided a forum for commentary and discussion. We hope that this volume will further the use of formal methods in clarifying one of the central problems of philosophy: that of our free human agency and its place in our indeterministic world.

2 State of the Art: BT, BST, *stit*, and CIFOL

In order to provide some background, we first give a brief and admittedly biased sketch of the current state of development of three formal frameworks that figure prominently in Nuel Belnap’s work on indeterminism and free action: the simple branching histories framework known as “branching time” (BT; Sect. 2.1), its relativistic spatio-temporal extension, branching space-times (BST; Sect. 2.2), and the “seeing to it that” (*stit*) logic of agency (Sect. 2.3). In Sect. 2.4, we additionally introduce case-intensional first order logic (CIFOL), a general intensional logic offering

² Readers interested in the concrete history can find some details in Appendix B at the end of this introduction.

resources for a first-order extension of the mentioned frameworks. CIFOL is a recent research focus of Belnap's, as reflected in his own contribution to this volume.

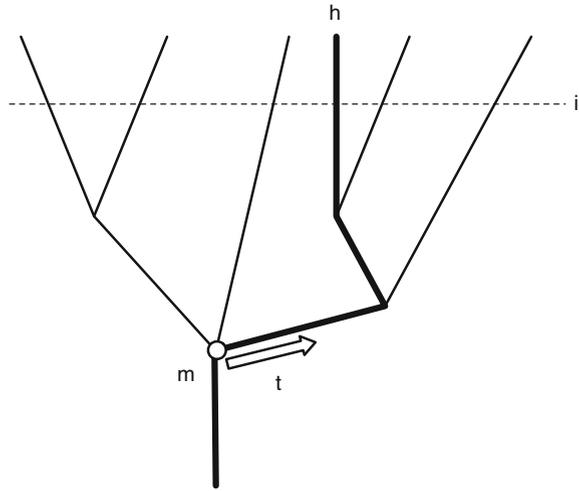
2.1 Branching Time (BT)

It is a perennial question of philosophy whether the future is open, what that question means, and what a positive or a negative answer to it would signify for us. The question has arisen in many different contexts—in science, metaphysics, theology, philosophy of language, philosophy of science, and in logic. The logical issue is not so much to provide an answer to the question about the openness of the future, nor primarily about its meaning and significance, but about the proper formal modeling of an open future: How can time and possibility be represented in a unified way? Thus clarified, the logical question of the open future is first and foremost one of providing a useful formal framework within which the philosophical issue of multiple future possibilities can be discussed.

In the light of twentieth century developments in modal and temporal logic, that logical question is one about a specific kind of possibility arising out of the interaction of time and modality. That kind of possibility may be called historical possibility or, in the terminology that Belnap favors, *real possibility*. A formal framework for real possibility must combine in a unified way a representation of past and future, as in temporal logic (tense logic), and of possibility and necessity, as in modal logic. That combination is not just interesting from a logical point of view—it is also of broader philosophical significance. To mention one salient example, the interaction of time and modality reflects the loss of possibilities over time that seems central to our commonsense idea of agency.

Working on his project of tense logic, Arthur Prior devoted his first book-length study to the topic of *Time and modality* (Prior 1957). A leading idea was that temporal possibility should somehow be grounded in truth at some future time, where time is depicted as linearly ordered. In 1958, Saul Kripke suggested a different formal framework, making use of partial orderings of moments. His exchange with Prior is documented in Ploug and Øhrstrøm (2012). The leading idea, which Prior took up and developed in his later book, *Past, present and future* (Prior 1967), was that the openness of the future should be modeled via a tree of histories (or chronicles) branching into the future. In terms of the partial ordering of moments m , a history h is a maximal chain (a maximal linearly ordered subset) in the ordering—graphically, one complete branch of the tree, representing a complete possible course of events from the beginning till the end of time (see Fig. 1). If the future is not open, all possible moments are linearly ordered, and there is just one history; if the future is open, however, the possible moments form a partial ordering in which there are multiple histories. In that case, we can say that there are incompatible possibilities for the same clock time (or for the same instant, i), which lie on different histories. Tomorrow, as Aristotle's famous example goes, there could be a sea-battle, or there could be none, and nothing yet decides between these two future possibilities.

Fig. 1 BT structure. m is a moment, and h , indicated by the bold line, is one of the structure's six histories. t is one of the three distinct transitions originating from m . The dashed line, i , indicates an instant, a set of moments at the same clock time in different histories. The future direction is up



This approach to modeling indeterminism has come to be known as “branching time” (BT), even though Belnap rejects the label on the ground that time itself “never [...] ever “branches”” (Belnap et al. 2001, 29). It is indeed better to speak of branching histories, since it is the histories that branch off from each other at moments. The label “branching time” is, however, well entrenched in the literature. Prior’s own development of BT was not fully satisfactory, but Thomason (1970) clarified its formal aspects in a useful way, adding even more detail in his influential handbook article on “Combinations of tense on modality” (Thomason 1984). The most versatile semantic framework for BT, which goes under Prior’s heading of “Ockhamism” due to an association with an idea of Ockham’s, posits a formal language with temporal operators (“it was the case that”, “it will be the case that”) and a sentential operator representing real possibility. The semantics of these operators is given via BT structures. The distinctive mark of Ockhamism is that it takes the truth of a sentence about the future to rely on (minimally) two parameters of truth, a temporal moment and a history containing that moment.³

The Ockhamist set-up can be developed in various ways, and Belnap has explored many of these in detail. We mention a number of salient issues and give a few references. The contributions to this volume by Brown and by Garson both develop further foundational issues of BT: Brown relates, *inter alia*, to the notion of a possible world that can ground alethic modalities; Garson connects the issue of the open future to the question of what is expressed by the rules of propositional logic and argues for a natural open future semantics that allows one to rebut logical arguments for fatalism.

³ See the article by Peter Øhrstrøm in this volume for discussion and historical details, including a hypothetical response to Belnap’s employment of the BT framework on Ockham’s behalf.

- BT, and also the earlier system of tense logic, brings out the dependence of the truth of a sentence on a suite of *parameters of truth*. For a simple temporal language, the truth of a sentence such as “Socrates is sitting” depends only on the moment with respect to which the sentence is evaluated. In Ockhamism, a sentence expressing a future contingent, such as “Socrates will be sitting at noon”, or indeed “There will be a sea-battle tomorrow”, is true or false relative to (minimally) two parameters, a moment and a history. Such a sentence, evaluated at some moment, can be true relative to one history and false relative to another. Relativity of truth to parameters of truth is nothing new or uncommon—it occurs already in standard predicate logic (see the next point). But in Ockhamism, one is forced to consider the issue of parameters of truth explicitly and in detail. A recognition of that issue has paved the way for a general semantics for indexical expressions (also known as “two-dimensional semantics”), as in the work of Kamp (1971) and Kaplan (1989). Belnap has pointed out the far-reaching analogy between “modal” parameters (such as m and h in Ockhamism) and an ordinary assignment of values to variables in predicate logic (Belnap et al. 2001, Chap. 6B).
- Working with this analogy, there is the interesting issue of how, given a context of utterance (or more generally, a context of use), parameters of truth receive a value that can be used in order to assign truth values to sentences. Belnap et al. (2001, 148f.) discuss this under the heading of “stand-alone sentences”; MacFarlane (2003) speaks of the issue of “postsemantics”. In the case of the variables in predicate logic, it seems quite clear that unless some value has been assigned to x , the sentence “ x is blue” cannot have a truth value. If all we have is “ x is blue”, the best we can do is prefix a quantifier, e.g., to read such a sentence as universally quantified, “for all x , x is blue”. In Ockhamism, a sentence minimally needs *two* parameters, a moment m and a history h containing m , in order to be given a truth value. How do these parameters receive a value? It seems plausible to assume that a context of utterance provides a moment of the context that can be used as an initial value for m . But what about h ? We make assertions about the future, but in an indeterministic partial ordering, there will normally be many different histories containing the moment m ; there is no unique “history of the context” to give the parameter h its needed value. This problem is known as the *assertion problem*. It does not seem that quantification provides a way out. Universal quantification in the semantics (an option known under Prior’s term “Peirceanism”) seems out of the question—when we say that it is going to rain tomorrow, we are not saying that it will necessarily rain tomorrow, i.e., that it will rain on *all* histories containing the present moment. When it turns out to be raining on the next day, we are satisfied and say that our assertion was true when made; we do not retract it when we are informed that sunshine was really possible (even though it didn’t manifest). These considerations also speak against the option of quantifying over the history parameter outside of the recursive semantics (“postsemantically”), as in supervaluationism (Thomason 1970). Similarly, one argues against existential quantifications over the relevant histories on the ground that when we say that it will be raining, we are claiming more than that it is possible that it will be raining. (On that option, we would have to say that both “It will be raining tomorrow” and

“It will not be raining tomorrow” are true, which sounds contradictory.) So, how do we understand assertions about the future?

- Together with Mitchell Green, Belnap has given a forceful statement of the problem of the uninitialized history parameter in Ockhamism and argued that it needs to be met head-on. According to Belnap and Green (1994), it will not do to posit a representation of “the real future” as a metaphorical “thin red line” singling out one future possibility above all others. They argue that marking any history as special, or real, would mean to deny indeterminism. (So, do not mistake the boldface line marking history h in Fig. 1 as indicating any special status for that history.) A number of solutions to the assertion problem have been discussed in the literature. Belnap (2002a) has argued that we can employ a second temporal reference point in order to assess future contingents later on. Before they can be assessed, a speech-act theoretic analysis can show their normative consequences. Here Belnap relies on the theory of word-giving developed by Thomson (1990). The current state of the debate appears to be that a “thin red line” theory is a consistent option from a logical point of view, but disagreements over the metaphysical pros and cons remain. In this volume, Øhrstrøm’s contribution gives a well-argued update on this discussion and its historical predecessors, while Green holds that a “thin red line” comes at an unnecessarily high metaphysical cost and argues that a speech-act theoretic understanding of our assertion practices is also possible.⁴
- Belnap has pointed out the importance of the notion of immediate, “local” possibilities for the proper understanding of the interrelation of time and modality. He finds in von Wright (1963) the notion of a “transition”, which is formally analyzed to be an initial paired with an immediately following outcome (Belnap 1999). Given an initial moment in a branching tree of histories, such a transition singles out a bundle of histories all of which remain undivided for at least some stretch of time. (Technically, one uses the fact that the relation of being undivided at a moment m is an equivalence relation on the set of histories containing m .) In Fig. 1, “ t ” indicates one of the three transitions (bundles of histories) branching off at m . Histories can then be viewed as maximal consistent sets of transitions. This allows for a generalization of the Ockhamist framework: instead of taking the parameters of truth to involve a moment/history pair m/h , one can employ a moment/set-of-transitions pair, m/T . Since sets of transitions are more fine-grained than whole histories, they can be used to represent the relative contingency of statements about the future, extending MacFarlane’s notion of a “context of assessment”. See Müller (2013a) and Rumberg and Müller (2013) for some preliminary results on this approach.
- Unlike theories developed in computer science, BT does not come with the assumption that the partial ordering of moments be discrete. While this assumption is certainly appropriate for many applications, it would trivialize some issues that can be usefully discussed in BT. An important case in point is the topology of branch-

⁴ For a recent defense of the “thin red line”, see also Malpass and Wawer (2012). MacFarlane (2014), on the other hand, defends assessment-relative truth of future contingents via his postsemantic approach.

ing. Assume that there are two continuous histories branching at some moment: Is there a last moment at which these two histories are undivided (a “choice point”), with the alternatives starting immediately afterwards, or should there be two alternative first moments of difference between these histories, so that there is no last moment of undividedness? McCall (1990) has illustrated these topologically different options. In BT, while assuming the existence of choice points is sometimes technically convenient, it makes no important difference which way one decides, as there is an immediate transformation of one representation into the other. This situation changes remarkably once we move to branching space-times.

2.2 Branching Space-Times (BST)

Branching space-times (BST) is a natural extension of the branching time framework, retaining the idea of branching histories for representing indeterminism but adding a formal representation of space in a way that is compatible with relativity theory. Belnap (2012) motivates the development of his theory of BST (Belnap 1992), which we will call BST1992, in the following way: Start with Newtonian space-time, which has an absolute (non-relativistic) time ordering and is deterministic. One way to modify this theory is to allow for indeterminism while sticking to absolute time. This corresponds to BT, in which the moments are momentary super-events stretching all of space. Another way to modify Newtonian space-time is to move to relativity theory, in which the notion of absolute simultaneity is abandoned in favor of a notion of simultaneity that is relative to a frame of reference. Combining the two moves, one arrives at a theory that is indeterministic (like BT) *and* relativistic (like relativistic space-time). Histories are no longer linear chains of moments ordered by absolute time, but whole space-times. Correspondingly, branching occurs not at space-spanning moments, but locally, at single possible point events.

The main technical innovation that makes BST1992 work, is the definition of a history not as a linear chain, but as an upward directed set in a partial ordering: a history contains, for any two of its members, a possible point event such that the two given members are in its causal past. In this way, one can work out branching history structures whose individual histories are all, e.g., Minkowski space-times (Müller 2002; Wroński and Placek 2009; Placek and Wroński 2009).

Historically, the origins of BST are somewhat different from the pedagogical set-up chosen by Belnap (2012). The story is interesting because it testifies to the mentioned intimate interrelation between indeterminism and agency. In the *stit* (“seeing to it that”) approach to the logic of agency, the truth conditions for “agent α sees to it that ϕ ” invoke the Ockhamist (BT-)parameters m/h . Briefly, for such a sentence to be true relative to moment m and history h , the agent α has to guarantee the outcome ϕ , which must not otherwise be guaranteed at m , by a choice determined by h . (See Sect. 2.3 for details.) Clearly, a single agent framework can only be the start; in fact, *stit* catered for multiple agents from the beginning. Now, intuitively speaking, what agents α and β choose to do at any given moment, should be independent: everybody

makes their own choices. It is reasonable to assume that this independence is guaranteed if agents α and β make their choices at different places at the same time, which implies that these choices are causally independent. But in a BT-based framework, there is no direct way to model that spatial separation. The solution in BT-based *stit* is, therefore, to introduce an additional axiom demanding independence. (See the contribution to this volume by Marek Sergot for a critical discussion of that axiom.) It would be much nicer if the agents' locations were modeled internally to the formalism, and the independence of their choices could accordingly be attributed to their spatial separation. An adequate notion of space-like relatedness is available in relativity theory, starting with Einstein's special theory of 1905. BST allows for a clear definition of space-like relatedness based on the underlying partial ordering: Two possible point events e and f are space-like related iff they are not order-related, but have a common upper bound (which guarantees that there is a history—a possible complete course of events—to which they both belong). Once agents are incorporated in BST (idealized as pointlike to begin with; see Belnap (2005a, 2011)), their choices can be taken to be events on their world-lines, and causal independence of such events can be directly expressed via space-like relatedness.

One can thus see two relevant motivations for constructing BST: as a relativistic extension of BT, and as a natural background theory for multi-agent logics of agency. The resulting quest for a reasonable framework for BST was mostly one of finding a useful definition of a history, and of fixing a number of topological issues, which become crucial in this development. Based on considerations of the causal attribution of indeterministic happenings, Belnap (1992) opts for the so-called “prior choice postulate”, which guarantees the existence of choice-points: For anything that happens in one history rather than in another, there is some possible point event in the past that is shared among the two histories in question, and which is maximal in their intersection. This postulate, together with continuity requirements, fixes to a large extent the topological structure of BST 1992.⁵ Figure 2 depicts a BST1992 structure with four histories, each of which is isomorphic to Minkowski space-time.

As in the case of BT, we mention a number of important issues and developments in BST to which Belnap has contributed. It will be obvious that he has been of central importance to all of them.

- To begin with topology, the original paper (Belnap 1992) mentions an approach to defining a topology for BST1992 that brings together different ideas from the theory of partial orders and from relativity theory. This topology, which Belnap attributes to Paul Bartha, has been researched in recent work by Placek and Belnap (2012); see also the contribution to this volume by Tomasz Placek. Naturally, the topological structure of a model of BST1992, which incorporates many incompatible histories, turns out to be non-Hausdorff (containing inseparable points);

⁵ There are related frameworks for incorporating space-time and indeterminism. An early description occurs in Penrose (1979); see also the references in Müller (2011a). McCall (1994) gives an informal description of branching models incorporating a spatial aspect; Strobach (2007) discusses alternatives in space-time from the point of view of defining logical operators. See also the remarks on topology and on general relativity's challenges for BST in the main text below.

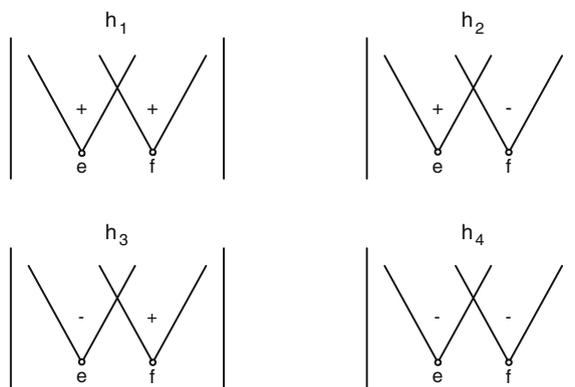


Fig. 2 Schematic diagram of a BST structure. e and f are choice points with two outcomes each, schematically denoted “+” and “-”. The four histories h_1, \dots, h_4 overlap outside the W-shaped forward lightcones of the choice points and in those parts of the light cones above e and f for which the labels coincide. The choice points e and f belong to all four histories. As in BT diagrams, the future direction is up

a single history is however typically Hausdorff. This makes good sense given indeterminism: If different possibilities exist for the same position in space-time, the corresponding possible point events may be topologically inseparable in the full indeterministic model.

- These topological observations are linked to the question whether BST can be viewed as a space-time theory. Earman (2008) has asked a pointed question about the tenability of BST as a space-time theory, sharply criticizing McCall's (1994) version of BST and raising doubts about Belnap's framework. His main challenge is to clarify the meaning of non-Hausdorffness that occurs in BST, since in space-time theories this is a highly unwelcome feature. Some recent literature, including Tomasz Placek's contribution to this volume, has clarified the situation considerably, highlighting the difference between branching *within* a space-time, which indeed has unwelcome effects well known to general relativists, and the BST notion of branching histories, in which the histories are individually non-branching space-times. The connection between BST and general relativity is only beginning to be made, and a revision of Belnap's prior choice principle may be in order to move the two theories closer to each other. (Technically, the issue is that the prior choice principle typically leads to a violation of local Euclidicity, which is, however, presupposed even for generalized, non-Hausdorff manifolds.) Apart from Placek's contribution, see also Sect. 6 of the contribution by Pleitz and Strobach, and Müller (2013b).
- Another area of physics that may be able to interact fruitfully with the BST framework is quantum mechanics. As BST incorporates both indeterminism and space-like separation, it seems to be especially well suited for clarifying the issue of space-like correlations in multi-particle quantum systems, pointed out in a famous

paper by Einstein et al. (1935). Following some pertinent remarks already in the initial paper by Belnap (1992), there have been some applications of the BST framework in this area, starting with Szabó and Belnap (1996), who target the three-particle, non-probabilistic case of Greenberger-Horne-Zeilinger (GHZ) states. These modeling efforts are connected with research on various types of common cause principles—see Hofer-Szabó et al. (2013). Placek (2010) brings into focus the epistemic nature of observed surface correlations vis-à-vis an underlying branching structure. For some remarks on a link between BT- or BST-like branching history structures and the quantum-mechanical formalism of so-called consistent histories, see Müller (2007).

- Even independently of applications to quantum physics, which may help to show the empirical relevance of the BST framework, there is the structural issue of how space-like correlations can be modeled in BST. Corresponding formal investigations were begun by Belnap (2002b) and continued in Belnap (2003), where the equivalence of four different definitions of modal correlations in BST1992 is proved. The basic observation is that it is possible to construct BST models in which the local possibilities at space-like separated choice points do not always combine to form global possibilities, i.e., histories. The simplest case corresponds exactly to the phenomenon pointed out in Einstein et al. (1935): Given a certain two-particle system, once its components are separated spatially, certain measurement outcomes for the components are perfectly correlated, meaning that it is impossible that a specific outcome on one side is paired with a specific outcome on the other side, even though no single outcome on either side is excluded. For an illustration, think of Fig. 2 with histories h_2 and h_3 missing: both choice points e and f then have two possible outcomes each, but the respective outcomes are perfectly modally correlated, admitting only joint outcomes $++$ and $--$. Müller et al. (2008) generalize Belnap's mentioned BST1992 results to incorporate cases of infinitely many correlated choice points. In this generalization, the notion of a transition, mentioned above in connection with BT, is crucial. For the use of sets of transitions to describe possibilities in BST, see also Müller (2010).
- The idea of (sets of) transitions as representatives of local possibilities is also the driving motor behind Belnap's highly original analysis of indeterministic causation (Belnap 2005b). In his approach, the relata of a singular causal statement “ C caused E ” are a transition (E) and a set of (basic) transitions (C). For a given effect E , described as “initial I followed by outcome O ”, it is possible in BST1992 to single out the relevant choice points (past causal loci) of that transition, and to describe the cause in terms of basic transitions in the past of O that lead from a choice point to one of its immediate local possibilities. These *causae causantes*, as Belnap calls them, are themselves basic causal constituents of our indeterministic world. Using various generalizations of the notion of an outcome, Belnap can prove that the *causae causantes* of an outcome constitute INUS conditions: insufficient but nonredundant parts of an unnecessary but sufficient condition for the occurrence of the outcome. (The notion of an INUS condition is famously from Mackie (1980).) Belnap's analysis provides a strong ontological reading of “causation as difference-

making” that appears to be well suited to modeling the kind of causation involved in human agency.

- Another useful employment of transitions in BST is in defining probabilities. Groundbreaking work was done by Weiner and Belnap (2006); a generalization to sets of transitions is given in Müller (2005), published earlier but written later. Paralleling earlier but independent work by Weiner, Müller (2005) shows that considerations of probability spaces lead to topological observations about BST as well. A general overview of probability theory in branching structures is given by Müller (2011b).

The basic idea of defining probability spaces in BST is to start with local probability spaces, defined on the algebra of outcomes of a single choice point. The interesting issue is how to combine such local probability spaces to form larger ones. Here it becomes crucial to consider consistent sets of transitions and to exclude pseudo-events whose probabilities make no sense. Müller (2005) offers the notion of a “causal probability space” in an analysis of which probability spaces can be sensibly defined in BST.

- The formal structure of BST is rich and multiply interpretable. This volume's contributions by Strobach and by Pleitz and Strobach testify to the versatility of the BST framework by providing a biological interpretation. Further developments are to be expected in the interaction between BST and the *stit* logic of agency.

2.3 Seeing to it That (*stit*)

We already remarked on some aspects of the *stit* framework that show its relation to branching histories frameworks and specifically to the development of BST. *Stit* logic is based on BT structures and uses the Ockhamist parameters of truth m and h , as introduced in Sect. 2.1. In order to represent agents and agency, BT structures are augmented via a set A of agents and an agent-indexed family of choices at moments, $Choice_m^\alpha$, which represent each agent's alternatives at each moment as a partition of the histories passing through that moment. These choices must be compatible with the local granularity of branching (the transition structure) resulting from the underlying BT structure: Agents cannot choose between histories before they divide in the structure (“no choice before its time”).

The semantic clause for “ α sees to it that ϕ ”, evaluated at m/h , has two parts: a positive condition, demanding that α must settle the truth of ϕ through her choice, and a negative condition, which excludes as agentive those ϕ whose truth is settled anyway. More specifically, there are two different developments of *stit*, which Belnap et al. (2001) call the “deliberative *stit*” (*dstit*) and the “achievement *stit*” (*astit*), respectively. The difference between them is one of perspective on what it is that the agent sees to. Both are built upon the mentioned BT structures with agents and their choices, but *astit* uses an additional resource, viz., a partitioning of the set of moments into so-called *instants* that mark the same clock time across different histories (Fig. 1 depicts one such instant, i). The book by Belnap, Perloff and

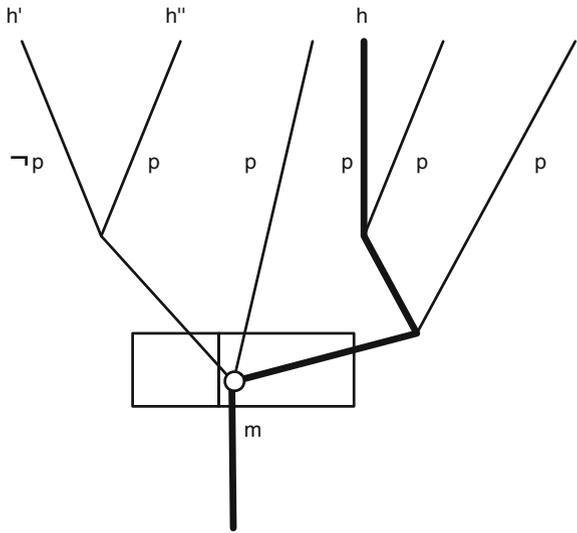
Xu, *Facing the future*, gives a comprehensive overview of a large number of developments in the *stit* framework, and is highly recommended as a general reference (Belnap et al. 2001). We leave many of the topics treated in that book, such as normative issues, strategies, word giving, or details of the resulting logics, to the side and describe just the basic frameworks, *astit* and *dstit*. Even though *astit* is historically earlier (Belnap and Perloff 1988), we start with a description of the simpler deliberative *stit*.—It is important to stress that while mentalistic notions such as deliberation are mentioned in the *stit* literature, the basic frameworks do not go beyond modeling the indeterministic background structure of agency; agents' beliefs and epistemic states do not play a role in the formal theory. This keeps the framework simple and general. Specific applications, however, can call for extra resources. The contributions to this volume by Bartha, Van Benthem and Pacuit, Broersen, Sergot, Vanderveken and Xu all testify to this: each discusses specific and useful additional details. Bartha adds utilities and probabilities in order to ground normative notions; Broersen also treats normative issues, via an Andersonian “violation” constant; Sergot models normativity via flagged (“red” or “green”) states. Van Benthem and Pacuit draw a comparison between *stit* and dynamic action logics, discussing a number of extensions that suggest themselves, including a dynamification of *stit*. Broersen adds probabilities for bringing about as well as subjective probabilities in order to anchor epistemic notions. Sergot employs a slightly different formal framework based on labeled transition structures, drops the independence of agents axiom, and emphasizes the importance of granularity of description for normative verdicts. Vanderveken adds a rich logic of propositional attitudes in order to analyze the logical form of proper intentional actions, extending the *stit* approach such as to give a logic of practical reason. Xu, in contrast, stays close to the austere *stit* framework; he explores in formal detail the extension of *stit* by group choices and group strategies. Further extensions of the basic *stit* approach are certainly possible.

Dstit was defined in Horty and Belnap (1995). The perspective is on securing a future happening due to a present choice, or deliberation. The positive clause for *dstit* demands that every history in the agent's current choice set satisfy the (future) outcome. The negative condition demands a corresponding witness for the violation of that outcome, which must belong to one of the other choices available to the agent. See Fig. 3 for an illustration; history h' fulfills the negative condition for $\alpha \text{ dstit} : p$, which is true at m/h . Large parts of *stit* can be developed without the negative condition, which greatly simplifies the logic; the corresponding *stit* operator is called *cstit*, after Chellas's employment of a similar idea in his analysis of imperatives (Chellas 1969). (A further simplification is possible if one assumes discrete time, see below.) Apart from the mentioned book by Belnap et al. (2001), see also Horty (2001).⁶

Belnap's historically first *stit* framework (Belnap and Perloff 1988) is based on the achievement *stit*, *astit*. As mentioned, instants are needed to define the *astit* operator. The perspective is different from that of *dstit*. For *astit*, a current result, or achievement, is attributed to an agent if there is a past witnessing moment at

⁶ For an independent, similar development, see also von Kutschera (1986).

Fig. 3 Illustration of *dstit*. The BT structure is that of Fig. 1. At moment m , the agent α has two possible choices, marked by the two boxes. (For the other moments, the choices are not indicated to avoid visual clutter.) On history h , but not on history h' nor on history h'' , α sees to it that p



which the agent's choice (as determined by the given history parameter) guaranteed the current result: All histories in that former moment's respective choice set must guarantee the result at the given instant (positive condition), and there must be another history passing the witnessing moment that does not lead to the result at that instant (negative condition). The logic of *astit* is interesting and quite complex; see Belnap et al. (2001, Chaps. 15–17).

In the recent literature, *dstit* plays the larger role. This may be due to its simpler logic, but perhaps also reflects the fact that the *dstit* operator is helpful for a formal representation of one of the main positions in the current free will debate, so-called *libertarianism*. According to the libertarian, free agency presupposes indeterminism. An influential argument given in favor of this assumption, Van Inwagen's so-called consequence argument (Van Inwagen 1983), proclaims that an action cannot be properly attributed to an agent if its outcome is already settled by events outside of the agent's control, and that would invariably be so under determinism. See the contribution to this volume by Robert Kane for a defense of libertarianism that points out the virtues of *stit* as a logical foundation for an intelligible account of free will based on indeterminism.

A helpful result in the logic of *dstit* is that refraining can itself be seen to be agentive, and that refraining from refraining amounts to doing. This result should be useful for clarifying the status of the assumption of alternative possibilities that is widely discussed in the free will debate and on whose merits or demerits much ink has been spilt. In *dstit*, if α sees to it that ϕ relative to the (Ockhamist) parameters m/h , this implies that there is a history h' containing m on which ϕ turns out false—that is the gist of the negative condition. As this history must lie in one of the agent's choices other than the one corresponding to h (this is due to the choices forming a

partition of the histories through m , and the positive condition demanding the truth of ϕ on all histories choice-equivalent to h), on that alternative, the agent *sees to it that she is not seeing to it that ϕ* . After all, making the choice corresponding to h' , α is not seeing to it that ϕ (since on h' , ϕ turns out false), but there is a history, viz., h , on which she *does* see to it that ϕ . So, “ α sees to it that she is not seeing to it that ϕ ” is the *stit* analysis of refraining. You can check that in Fig. 3, at m on history h' , the agent refrains from future p (Fp) in exactly that sense. It is clear that the alternative of refraining from ϕ does *not* have to amount, on that analysis, to the agent’s possibly seeing to it that $\neg\phi$, even though this is often taken to be implied by the assumption of alternative possibilities. (In Fig. 3, there is no history on which α sees to it that $\neg Fp$.) In our view, *stit* provides some desperately needed clarity here.⁷ There is certainly much work to be done to integrate formal work on *stit* into the free will debate. See Kane’s contribution to this volume for a discussion of a number of additional steps towards a fuller account of indeterminism-based free will.

Outside of philosophy proper, *stit* has had, and continues to have, a significant influence on the modeling of agency in computer science and artificial intelligence. Many of the contributions to this volume testify to *stit*’s usefulness in this area. Usually, such applications of the framework give up the initial generality of BT models (which allow for continuous structures) in favor of discrete orderings. While this means a limitation of scope, it makes the framework much more tractable and thus, useful from an engineering point of view. The availability of a “next time” operator suggests that one can read a *dstit*- or *cstit*-like operator as “an agent secures an outcome at all choice-equivalent possible next moments”, thus doing away with a layer of complexity introduced by the usual handling of the future tense (which quantifies over all future moments on a given history, including moments that are far removed), and by the need for considering whole histories. In this volume, the contribution by Broersen explicitly builds upon discrete structures, and the transition system framework employed by Sergot is also typically discrete. Van Benthem and Pacuit in their contribution leave the basic *stit* framework unconstrained, but go on to employ the discrete view of stepwise execution that is basic for dynamic logic. With various refinements and extensions of *stit*, it seems fair to say that the computer science community currently provides the richest environment for the development of that framework. Interaction with the philosophical community can certainly prove to be beneficial for both sides, and we hope that this volume can be helpful in that respect.

It should also be stressed that while the *stit* framework has found many applications, it is by no means the only approach to the formal modeling of agency on the market. Two of the contributions to this volume draw explicit connections to other important existing frameworks. Sergot remains close to the *stit* framework,

⁷ We refrain from entering a lengthier discussion of the free will debate, which has turned into a maze of arguments, counterarguments and, not too infrequently, confusion and talking past each other. From among the recent original and helpful contributions to the debate, we mention Helen Steward’s plea for the libertarian position of “agency incompatibilism” (Steward 2012). She indicates connections to the *stit* framework as well (Steward 2012, 31).

but draws upon the formalism of Pörn (1977). Van Benthem and Pacuit provide a detailed comparison between the *stit* approach and the paradigm of dynamic logic that was developed in the formal study of computer programs. These comparisons are highly valuable, since they promise to help to bring related research lines operating in relative isolation closer together.

Since *stit* is so rich and multi-faceted, we do not attempt here to give an overview of recent developments akin to what we did for BT and BST above. We refer again to the book, *Facing the future* (Belnap et al. 2001), for the groundwork and a clear presentation of logical issues. For contemporary developments, we refer to the contributions in this volume.

2.4 Case-Intensional First Order Logic (CIFOL)

The development of all the three mentioned frameworks—BT, BST, and *stit*—is based on semantical considerations, though not necessarily with a view toward providing a semantics for an extant formal language. The common, semantically driven idea is to define structures that represent aspects of reality such that the truth or falsity of sentences can be discussed against the background of such a structure.

When one looks at applications that do relate to a formal language (such as the language of tense logic for BT), it turns out that most often, models based on the respective structures are thought of as providing the semantics for a *propositional* language, which does not use variables or quantifiers. This is probably mostly due to the fact that many actual applications arise in a computer science context, and propositional logic is computationally much more tractable than predicate logic. There may also still be a lingering worry about the tenability of quantified modal logic, even though Quine's influence is waning.⁸ But perhaps the main reason for the fact that there is not a lot of BT-based predicate logic (let alone a predicate logic based on BST, or on *stit*) is that it is hard to get it right. For philosophical purposes, it is, however, clear that we need to take individual things seriously—after all, we, the biological creatures populating this planet, are agents, and it is not always fruitful to reduce the representation of one of us to a mere label on a modal operator. Thus, one of the areas in which much further logical development is to be expected, is an adequate representation of things, their properties and their possibilities in an indeterministic setting.

Quantified modal logic (QML) has long been an area of interaction between logic and metaphysics, not always to the benefit of logic. One of the most interesting recent developments in Belnap's work on indeterminism and free action is connected with the attempt of developing a metaphysically neutral quantified modal logic, which would be driven by applicability rather than by underlying metaphysical assumptions. Consider the handling of variables. Most systems of QML assume that modal logic

⁸ For Quine's arguments against quantifying into modal contexts, see, e.g., Quine (1980, Chap. VIII). See Fine (2005, Chaps. 2 and 3) for extensive analysis and critique.

should be built on a modal parameter of truth that specifies a “possible world”. Also, typically, a variable functions as a rigid designator: Each possible world comes with its domain of individuals (the world’s “inhabitants”), and a variable designates the same individual in any world. Alternatively, a counterpart relation between the domains and a corresponding handling of variables is discussed.⁹ Both moves make a certain view of the metaphysical status of individuals part of the quantificational machinery of QML. Accordingly, such logics cannot be used to represent dissenting metaphysical views about individuals. It would seem, however, that one of the main virtues of using a logical formalism is that it provides an arena in which different views can be formulated and arguments in their favor or against them can be checked. What good is a quantified modal logic if it does not allow one to discuss different theories and arguments about the metaphysical status of individuals?

Belnap argues for a broader, more general approach to QML that is based on a neglected but useful framework for quantified modal logic developed in the interest of clarifying arguments arising in the empirical sciences. Aldo Bressan (1972) developed his case-intensional approach to QML out of his interest in the role of modality in physics. His system is higher order and includes a logicist construal of the mathematics necessary for applications in physics; this makes it highly complex and may have stood in the way of its wider recognition or application. Belnap (2006) provides a useful overview of the general system. For many purposes it is, however, sufficient to look at the first-order fragment of Bressan’s system, and to develop that as a stand-alone logical framework. One guiding idea is generality: instead of developing a modal logic based on the idea of a “possible world”, or a temporal logic that is geared towards truth at a time, it is better to work with a general notion of a modal parameter of truth that we may call a *case*. This accords with ordinary English usage, and justifies S5 modalities built upon cases: necessary is what is true in any case; something is possible if there is at least one case in which it is true. Another guiding idea is uniformity. Rather than following standard systems of QML, which treat variables, individual constants and definite descriptions in widely different ways, one can use the most general idea of a term with an extension in each case, and an individual intension that represents the pattern of variation of the extension across cases. (Technically, the intension is the function from cases to extensions, and the extension at a case is the intension-function applied to that case. This recipe is followed uniformly for all parts of speech, generalizing Carnap’s (1947) method of extension and intension.) Correspondingly, the most general option is used for predication as well: predication is not forced to be extensional, but is generally intensional, such that a one-place predicate for each case provides a function that maps intensions to truth values. This rich and uniform background provides for a simple yet powerful definition of sortal properties as allowing for the tracing of individuals from case to case. See Belnap and Müller (2013a) for a detailed description of the resulting framework of case-intensional first order logic (CIFOL). The framework has recently been extended to cases in a branching histories framework (Belnap and Müller 2013b). This application of CIFOL helps to dispell worries that have been raised against

⁹ For an in-depth overview, see Kracht and Kutz (2007).

the idea of individuals in branching histories, such as famously in Lewis's argument against branching (Lewis 1986, 206ff): Using the resources of CIFOL, it is possible to model individuals and sortal properties successfully in a branching histories framework. Good news, surely, for those of us who believe that we are just that: individual agents facing an open future of possibilities.

In line with the development of BT, BST and *stit*, CIFOL is developed from a semantical point of view. The interface with a formal logical language is, however, much more pronounced in the case of CIFOL—the fact that we are considering a predicate logic necessitates close attention to the syntax as well. (For example, as the framework is required to remain first-order, while lambda-abstraction is unfettered, lambda-predicates may only occur in predicate position.) Naturally, it is to be expected that there can be fruitful discussions of CIFOL's proof theory and metatheory. Nuel Belnap, in his contribution to this volume, gives a highly interesting overview of a truth theory that can be developed within CIFOL+, a minimal extension of CIFOL. Given the framework's intensionality, it is possible to define terms representing the cases, and based on those, one can develop the theory of the mixed nector “that Φ is true at case x ”. You will, we hope, not go wrong in expecting further striking results about CIFOL and its connection to indeterminism and free action in the near, albeit open future.

Acknowledgments Research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC Grant agreement nr 263227, and from the Dutch Organization for Scientific Research, grant nr NWO VIDI 276-20-013. I would like to thank Nuel Belnap for continuing inspiration, and both Nuel Belnap and Antje Rumberg for helpful comments on a previous draft.

Open Access This chapter is distributed under the terms of the Creative Commons Attribution Noncommercial License, which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

Appendix A: Abstracts of the Papers in this Volume

Paul Bartha (University of British Columbia): Decisions in Branching Time

This paper extends the deontic logic of Horty (*Agency and deontic logic*, 2001) in the direction of decision theory. Horty's deontic operator, the *dominance ought*, incorporates many concepts central to decision theory: acts, causal independence, utilities and dominance reasoning. The decision theory associated with dominance reasoning, however, is relatively weak. This paper suggests that deontic logic can usefully be viewed as *proto-decision theory*: it provides clear foundations and a logical framework for developing norms of decision of varying strength. Within Horty's framework, deontic operators stronger than the dominance ought are defined for decisions under ignorance, decisions under risk, and two-person zero-sum games.

Nuel Belnap (University of Pittsburgh): Internalizing case-relative truth in CIFOL+

CIFOL is defined in Belnap and Müller (*J Phil Logic* 2013) as the first-order fragment of Aldo Bressan’s higher-order modal typed calculus MC^v Bressan based his calculus on Carnap’s “method of extension and intension”: In CIFOL, truth is relative to “cases,” where cases play the formal role of “worlds” (but with less pretension). CIFOL+ results by following Bressan in adding term-constants \mathbf{t} for the true and \mathbf{f} for the false, and a single predicate constant, P_0 , which together with a couple of simple axioms enable the representation of “sentence Φ is true in case x ” by means of a defined expression, $T(\Phi, x)$, where Φ is the sentence of CIFOL+ in question and where x ranges over a defined family of “elementary cases.” (Whereas being a *case* is defined in the semantic metalanguage, *elementary cases* are squarely in the (first order) domain of CIFOL+.) A suitable suite of axioms guarantees that one can prove (in CIFOL+) that there is exactly one elementary case, x , such that x happens (i.e., such that $x = \mathbf{t}$), a fact that underlies the equivalence of $\Box(x = \mathbf{t} \rightarrow \Phi)$ and $\Diamond(x = \mathbf{t} \wedge \Phi)$. (Proofs are surprisingly intricate for first order modal logic). One can then go on to show that $T(\Phi, x)$ is well-behaved in terms of its relation to the connectives of CIFOL+, a result required for ensuring that $T(\Phi, x)$ is properly read as “ Φ is true in elementary case x .”

Jan Broersen (Utrecht University): A stit Logic Analysis of Morally Lucky and Legally Lucky Action Outcomes

Moral luck is the phenomenon that agents are not always held accountable for performance of a choice that under normal circumstances is likely to result in a state that is considered bad, but where due to some unexpected interaction the bad outcome does not obtain. We can also speak of moral misfortune in the mirror situation where an agent chooses the good thing but the outcome is bad. This paper studies formalizations of moral and legal luck (and moral and legal misfortune). The three ingredients essential to modelling luck of these two different kinds are (1) indeterminacy of action effects, (2) determination on the part of the acting agent, (3) the possibility of evaluation of acts and/or their outcomes relative to a normative moral or legal code. The first, indeterminacy of action, is modelled by extending stit logic by allowing choices to have a probabilistic effect. The second, deliberateness of action, is modelled by (a) endowing stit operators with the possibility to specify a lower bound on the change of success, and (b) by introducing the notion of attempt as a maximisation of the probability of success. The third, evaluation relative to a moral or legal code, is modelled using Andersons reduction of normative truth to logical truth. The conclusion will be that the problems embodied by the phenomenon of moral luck may be introduced by confusing it with legal luck. Formalizations of both forms are given.

Mark A. Brown (Syracuse University): Worlds Enough, and Time—Musings on Foundations

Belnap’s work on stit theory employs an Ockhamist theory of branching time, in which the fundamental possibilities within models are commonly taken to be moments of time, connected into a tree-like branching structure. In the semantics for alethic

modal logic, necessity is characterized by quantification over relevant possible worlds within a model, yet Belnap refers to an entire model of branching time as our world, seemingly leaving no room for non-trivial quantification over worlds within a single model.

This paper explores the question how the notion of possible worlds should be understood in relation to an Ockhamist framework, in order to be able to combine an account of alethic modalities with an account of branching time and stit theory. The advantages and drawbacks of several alternative approaches are examined.

James W. Garson (University of Houston): Open Futures in the Foundations of Propositional Logic

This paper weaves together two themes in the work of Nuel Belnap. The earlier theme was to propose conditions (such as conservativity and uniqueness) under which logical rules determine the meanings of the connectives they regulate. The later theme was the employment of semantics for the open future in the foundations of logics of agency. This paper shows that on the reasonable criterion for fixing meaning of a connective by its rule governed deductive behavior, the natural deduction rules for classical propositional logic do not fix the interpretation embodied in the standard truth tables, but instead express an open future semantics related to Kripke's possible worlds semantics for intuitionistic logic, called natural semantics. The basis for this connection has already been published, but this paper reports new results on disjunction, and explores the relationships between natural semantics and supervaluations. A possible complaint against natural semantics is that its models may disobey the requirement that there be no branching in the past. It is shown, however, that the condition may be met by using a plausible reindividuation of temporal moments. The paper also explains how natural semantics may be used to locate what is wrong with fatalistic arguments that purport to close the door on an open future. The upshot is that the open future is not just essential to our idea of agency, it is already built right into the foundations of classical logic.

Mitchell Green (University of Virginia): On Saying What Will Be

In the face of ontic (as opposed to epistemic) openness of the future, must there be exactly one continuation of the present that is what *will* happen? This essay argues that an affirmative answer, known as the doctrine of the Thin Red Line, is likely coherent but ontologically profligate in contrast to an Open Future doctrine that does not privilege any one future over others that are ontologically possible. In support of this claim I show how thought and talk about "the future" can be shown intelligible from an Open Future perspective. In so doing I elaborate on the relation of speech act theory and the "scorekeeping model" of conversation, and argue as well that the Open Future perspective is neutral on the doctrine of modal realism.

Robert Kane (University of Texas at Austin): The Intelligibility Question For Free Will—Agency, Choice And Branching Time

In their important work, *Facing the Future* (Oxford 2001), Nuel Belnap and his collaborators, Michael Perloff and Ming Xu, say the following (p. 204): “We agree with Kane [1996] that ... the question whether a kind of freedom that requires indeterminism can be made intelligible deserves ... our most serious attention, and indeed we intend that this book contribute to what Kane calls ‘the intelligibility question.’” I believe their book does contribute significantly to what I have called “the Intelligibility Question” for free will (which as I understand it is the question of how one might make intelligible a free will requiring indeterminism without reducing such a free will to either mere chance or to mystery and how one might reconcile such a free will with a modern scientific understanding of the cosmos and human beings). The theory of agency and choice in branching time that Belnap has pioneered and which is developed in detail in *Facing the Future* is just what is needed in my view as a logical foundation for an intelligible account of a free will requiring indeterminism, which is usually called libertarian free will. In the first two sections of this article, I explain why I think this to be the case. But the logical framework which Belnap et al. provide, though it is necessary for an intelligible account of an indeterminist or libertarian free will, is nonetheless not sufficient for such an account. In the remaining sections of the article (3–5), I then discuss what further conditions may be needed to fully address “the Intelligibility Question” for free will and I show how I have attempted to meet these further conditions in my own theory of free will, developed over the past four decades.

Peter Øhrstrøm (Aalborg University): What William of Ockham and Luis de Molina would have said to Nuel Belnap—A Discussion of some Arguments Against “The Thin Red Line”

According to A.N. Prior the use of temporal logic makes it possible to obtain a clear understanding of the consequences of accepting the ideas of indeterminism and free choice. Nuel Belnap is one of the most important writers who have contributed to the further exploration of these tense-logical ideas as seen in the tradition after Prior.

In some of his early papers Prior suggested the idea of the true future. Obviously, this idea corresponds to an important notion defended by classical writers such as William of Ockham and Luis de Molina.

Belnap and others have considered this traditional idea introducing the term, “the thin red line” (TRL), arguing that this idea is rather problematic. In this paper I argue that it is possible to respond to the challenges from Belnap and others in a reasonable manner. It is demonstrated that it is in fact possible to establish a consistent TRL theory. In fact, it turns out that there several such theories which may all be said to support the classical idea of a true future defended by Ockham and Molina.

Tomasz Placek (Jagiellonian University, Kraków): Branching for general relativists

The paper develops a theory of branching spatiotemporal histories that accommodates indeterminism and the insights of general relativity. A model of this theory can be viewed as a collection of overlapping histories, where histories are defined

as maximal consistent subsets of the model's base set. Subsequently, generalized (non-Hausdorff) manifolds are constructed on the theory's models, and the manifold topology is introduced. The set of histories in a model turns out to be identical with the set of maximal subsets of the model's base set with respect to being Hausdorff and downward closed (in the manifold topology). Further postulates ensure that the topology is connected, locally Euclidean, and satisfies the countable sub-cover condition.

Marek Sergot (Imperial College): Some examples formulated in a 'seeing to it that' logic—Illustrations, observations, problems

The paper presents a series of small examples and discusses how they might be formulated in a 'seeing to it that' logic. The aim is to identify some of the strengths and weaknesses of this approach to the treatment of action. The examples have a very simple temporal structure. An element of indeterminism is introduced by uncertainty in the environment and by the actions of other agents. The formalism chosen combines a logic of agency with a transition-based account of action: the semantical framework is a labelled transition system extended with a component that picks out the contribution of a particular agent in a given transition. Although this is not a species of the *stit* logics associated with Nuel Belnap and colleagues, it does have many features in common. Most of the points that arise apply equally to *stit* logics. They are, in summary: whether explicit names for actions can be avoided, the need for weaker forms of responsibility or 'bringing it about' than are captured by *stit* and similar logics, some common patterns in which one agent's actions constrain or determine the actions of another, and some comments on the effects that level of detail, or 'granularity', of a representation can have on the properties we wish to examine.

Niko Strobach (Westfälische Wilhelms-Universität Münster): In Retrospect. Can BST models be reinterpreted for what decisions, speciation events and ontogeny might have in common?

This paper addresses two interrelated topics: (1) a formal theory of biological ancestry (FTA); (2) ontological retrospect. The point of departure is a reinterpretation of Nuel Belnap's work on branching spacetime (BST) in terms of biological ancestry. Thus, Belnap's prior choice principle reappears as a principle of the genealogical unity of all life. While the modal dimension of BST gets lost under reinterpretation, a modal dimension is added again in the course of defining an indeterministic FTA where possible worlds are alternatives in terms of offspring. Indeterministic FTA allows to model important aspects of ontological retrospect. Not only is ontological retrospect a plausible account for the perspectival character of Thomason-style supervaluations, but it is shown to be a pervasive ontological feature of a world in development, since it is relevant for cases as diverse as speciation, the individual ontogeny of organisms and decisions of agents. One consequence of an indeterministic FTA which includes the idea of retrospect is that, contrary to what Kripke famously claims, species membership is not always an essential feature, but may depend on the way

the world develops. The paper is followed by a postscript by Martin Pleitz and Niko Strobach which provides a version of indeterministic FTA that is technically even closer to Belnap's BST than the one in this paper and which allows for a discussion of further philosophical details.

Martin Pleitz and Niko Strobach (Westfälische Wilhelms-Universität Münster): A Theory of Possible Ancestry in the Style of Nuel Belnap's Branching Space-Time

We present a general theory of possible ancestry that is a case of modal ersatzism because we do not take possibilities in terms of offspring as given, but construct them from objects of another kind. Our construction resembles Nuel Belnap's theory of branching space-time insofar we also carve all possibilities from a single pre-existing structure. According to the basic theory of possible ancestry, there is a discrete partially ordered set called a structure of possible worlds, any subset of which is called admissible iff it is downward closed under the ordering relation. A structure of possible worlds is meant to model possible living beings standing in the relation of possible ancestry, and the admissible sets are meant to model possible scenarios. Thus the Kripkean intuition of the necessity of (ancestral) origin is incorporated at the very core of our theory. In order to obtain a more general formulation of our theory which allows numerous specifications that might be useful in concrete biological modeling, we single out two places in our framework where further requirements can be implemented: Global requirements will put further constraints on the ordering relation; local requirements will put further constraints on admissibility. To make our theory applicable in an indeterminist world, we use admissible sets to construct the (possible) moments and (possible) histories of a branching time structure. We then show how the problem of ontological competition can be solved by adding an incompatibility partition to a structure of possible worlds, and conclude with some remarks about how this addition might provide a clue for developing a variant of the theory of branching space-time that can account for the trousers worlds of general relativity.

Johan van Benthem and Eric Pacuit (University of Amsterdam and University of Maryland at College Park): Connecting Logics of Choice and Change

This paper is an attempt at clarifying the current scene of sometimes competing action logics, looking for compatibilities and convergences. Current paradigms for deliberate action fall into two broad families: dynamic logics of events, and STIT logics of achieving specified effects. We compare the two frameworks, and show how they can be related technically by embedding basic STIT into a modal logic of matrix games. Amongst various things, this analysis shows how the attractive principle of independence of agents' actions in STIT might actually be a source of high complexity in the total action logic. Our main point, however, is the compatibility of dynamic logics with explicit events and STIT logics based on a notion that we call 'control'—and we present a new system of dynamic-epistemic logic with control that has both. Finally, we discuss how dynamic logic and STIT face similar issues when including further crucial aspects of agency such as knowledge, preference, strategic behavior, and explicit acts of choice and deliberation.

Daniel Vanderveken (University of Quebec at Trois-Rivières): Intentionality and minimal rationality in the logic of action

Philosophers have overall studied intentional actions that agents attempt to perform in the world. However the pioneers of the logic of action, Belnap and Perloff, and their followers have tended to neglect the intentionality proper to human action. My primary goal is to formulate here a more general logic of action where intentional actions are primary as in contemporary philosophy of mind. In my view, any action that an agent performs involuntarily could in principle be intentional. Moreover any involuntary action of an agent is an effect of intentional actions of that agent. However, not all unintended effects of intentional actions are the contents of unintentional actions, but only those that are historically contingent and that the agent could have attempted to perform. So many events which happen to us in our life are not really actions. My logic of action contains a theory of attempt, success and action generation. Human agents are or at least feel free to act. Moreover their actions are not determined. As Belnap pointed out, we need branching time and historic modalities in the logic of action in order to account for indeterminism and the freedom of action.

Propositions with the same truth conditions are identified in standard logic. However they are not the contents of the same attitudes of human agents. I will exploit the resources of a non classical predicative propositional logic which analyzes adequately the contents of attitudes. In order to explicate the nature of intentional actions one must deal with the beliefs, desires and intentions of agents. According to the current logical analysis of propositional attitudes based on Hintikka's epistemic logic, human agents are either perfectly rational or completely irrational. I will criticize Hintikka's approach and present a general logic of all cognitive and volitive propositional attitudes that accounts for the imperfect but minimal rationality of human agents. I will consider subjective as well as objective possibilities and explicate formally possession and satisfaction conditions of propositional attitudes. Contrary to Belnap, I will take into account the intentionality of human agents and explicate success as well as satisfaction conditions of attempts and the various forms of action generation. This paper is a contribution to the logic of practical reason. I will formulate at the end many fundamental laws of rationality in thought and action.

Ming Xu (Wuhan University): Group strategies and independence

We expand Belnap's general theory of strategies for individual agents to a theory of strategies for multiple agents and groups of agents, and propose a way of applying strategies to deal with future outcomes at the border of a strategy field. Based on this theory, we provide a preliminary analysis on distinguishability and independence, as a preparation for a general notion of dominance in the decision-theoretical approach to deontic logic.

Appendix B: On the History of *stit* and Branching Space-Times

Interview with Nuel Belnap, conducted at his home in Pittsburgh, March 15, 2013.
Interviewer: Thomas Müller.

TM: Let's talk about the origins of *stit*. Jan Broersen, one of our authors, mentioned that you had told him about the history one evening over dinner, when you were in Utrecht a couple of years ago. You developed some of that in seminars, in the 1980s?

NB: It started with a seminar I taught on Charles Hamblin's book, *Imperatives*, as far as I recall. Maybe two seminars, maybe just the one. I certainly worked out a good bit about *stit* for the seminar, writing out a few pages each week.

TM: Hamblin's book came out in 1987, with your preface, so this must have been the mid-1980s. The first *stit* paper came out in 1988, so that would fit temporally. Rich Thomason, whose work on branching histories theories for indeterminism forms part of the formal background for *stit*, was your colleague at the University of Pittsburgh until 1999, when he moved to the University of Michigan. You have often remarked that you were amazed by how long this theory was lying dormant, with the initial paper from 1970 and the *Handbook of philosophical logic* chapter published in 1984—there was virtually nothing happening in between. Thomason has some remarks on the deontic aspects of his approach.

NB: He did work out some deontic ideas, yes.

TM: For *stit* you were mainly working with Mickey Perloff, right? And then some graduate students were attracted as the project was building up momentum—for example, Jeff Horty, Mitch Green, and Ming Xu. What I find interesting is the interaction between the two projects, clarifying the foundations of indeterminism through the application of indeterministic models in the logic of agency, and building up the logic of agency against the background of branching histories models for indeterminism. Your book, *Facing the future*, exemplifies this nicely.

NB: The book must be right.—Mickey took part in the Hamblin seminar; we worked together for many years afterwards.

TM: The branching times framework—I assume you knew about that from much earlier? When Alan Anderson was at Manchester to work with Prior in the mid-60s, he would have brought back some ideas about that?

NB: Yes, I think so. Prior visited Alan in 1965 or so, he came to a dinner party at his house. That's when he had decided not to come to the U.S. any more, because of the Vietnam war.

TM: So the branching time framework was basically sitting there to be used, and you made the connection, not working on issues in branching time, but when thinking about how to model the content of an imperative?

NB: Hamblin's book is on imperatives, yes. There's a mini-history of approaches to the modal logic of agency early in the book, *Facing the future*.

TM: When you started working on *stit*, was that working out the theory of a single agent first, with other agents entering the theory only later?

NB: No, the multi-agent case was in there from the beginning. The other agents didn't do anything, to begin with.

TM: There is the "independence of agents" axiom in multi-agent *stit*: "Something happens"; no matter what one agent chooses at a moment, all other possible choices of the other agents must be compatible with that. That was the nucleus of the project of branching space-times, I think Paul Bartha told me about that at one point?

NB: I do remember that I had the main ideas of branching space-time in the late 1980s, and I was shopping them around. Every visitor to the department got an hour of that. That was before the paper was published in 1992.

TM: Chris Hitchcock told me that he was there "when it happened".

NB: That was a small seminar, I think Chris and Philip Kremer were the only students in the class.—I don't have any records on what and who I was teaching. I had seven four-drawer file cabinets at the department, and when I retired a few years ago I just asked the secretary, Connie, to get rid of them.

TM: How did the main ideas come about?

NB: I learned about directed sets from Dana Scott. Not when he was at Carnegie Mellon University in Pittsburgh, but long before then. We overlapped at Oxford in 1970. Directed sets is really what made branching space-times go, it's the basis for the definition of a history. That idea had been with me for many years.

TM: This is a recurring theme in our discussions: It takes time. Ideas can take 20 years, and then they reappear, or become salient all of a sudden.

NB: They cook a long time.

TM: For me it's now 15 years since I first read the paper on branching space-times—and there's still a lot for me to discover, like the one footnote on topology that has driven a small industry over the last couple of years.

NB: I was just rereading it earlier the day, in order to see whether I could find the right platform for the method of extension and intension that we are working on now. I didn't get anywhere, though.

TM: It's good that you made that postprint, ten years after the first publication. That shows some progress.

NB: In the branching space-times paper in the beginning I had a substantial section on agency, which I was persuaded to disavow.

TM: There is a gap of more than ten years between the 1992 publication of the BST paper and your published work on agency in BST, starting around 2005.

NB: The connection was there from the start.

TM: Thanks, Nuel.

References

- Belnap, N. 1992. Branching space-time. *Synthese* 92(3): 385–434 (see also the postprint 2003, available on philsci-archive, <http://philsci-archive.pitt.edu/1003/>).
- Belnap, N. 1999. Concrete transitions. In *Actions, Norms, Values: Discussions with Georg Henrik von Wright*, ed. G. Meggle, 227–236. Berlin: de Gruyter.
- Belnap, N. 2002a. Double time references: Speech-act reports as modalities in an indeterminist setting. In *Advances in modal logic*, vol. 3, eds. F. Wolter, H. Wansing, M. de Rijke and M. Zakharyashev, 37–58. Singapore: World Scientific.
- Belnap, N. 2002b. EPR-like “funny business” in the theory of branching space-times. In *Non-locality and modality*, ed. T. Placek and J. Butterfield, 293–315. Dordrecht: Kluwer.
- Belnap, N. 2003. No-common-cause EPR-like funny business in branching space-times. *Philosophical Studies* 114: 199–221.
- Belnap, N. 2005a. Agents and agency in branching space-times. In *Logic, thought and action*, ed. D. Vanderveken, 291–313. Berlin: Springer.
- Belnap, N. 2005b. 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. 2006. Bressan’s type-theoretical combination of quantification and modality. In *Modality Matters. Twenty-five Essays in Honor of Krister Segerberg*, eds. H. Lagerlund, S. Lindström, and R. Sliwinski, 31–53. Department of Philosophy, Uppsala University, Uppsala.
- Belnap, N. 2011. Prolegomenon to norms in branching space-times. *Journal of Applied Logic* 9(2): 83–94.
- Belnap, N. 2012. Newtonian determinism to branching space-times indeterminism in two moves. *Synthese* 188(1): 5–21.
- Belnap, N. and M. Green. 1994. Indeterminism and the thin red line. *Philosophical Perspectives* 8:365–388 (ed. J. Tomberlin).
- Belnap, N., and T. Müller. 2013a. CIFOL: Case-intensional first order logic (I). Toward a logic of sorts. *Journal of Philosophical Logic*. doi:[10.1007/s10992-012-9267-x](https://doi.org/10.1007/s10992-012-9267-x).
- Belnap, N. and T. Müller. 2013b. BH-CIFOL: Case-intensional first order logic. (II) Branching histories. *Journal of Philosophical Logic*. doi:[10.1007/s10992-013-9292-4](https://doi.org/10.1007/s10992-013-9292-4).
- Belnap, N., and M. Perloff. 1988. Seeing to it that: A canonical form for agentives. *Theoria* 54:167–190. Corrected version in Kyburg Jr., H.E., R.P. Loui and G.N. Carlson, eds., 1990. *Knowledge representation and defeasible reasoning*, *Studies in cognitive systems*, vol. 5, 167–190. Dordrecht: Kluwer.
- Belnap, N., M. Perloff, and M. Xu. 2001. *Facing the future. Agents and choices in our indeterminist world*. Oxford: Oxford University Press.
- Bressan, A. 1972. *A general interpreted modal calculus*. New Haven: Yale University Press. (Foreword by Nuel D. Belnap Jr.)
- Carnap, R. 1947. *Meaning and necessity: A study in semantics and modal logic*. Chicago: University of Chicago Press.
- Chellas, B.F. 1969. *The logical form of imperatives*. Stanford: Perry Lane Press.
- Earman, J. 2008. Pruning some branches from branching space-times. In *The Ontology of Spacetime II*, ed. D. Dieks, 187–206. Amsterdam: Elsevier.
- Einstein, A., B. Podolsky, and N. Rosen. 1935. Can quantum-mechanical description of physical reality be considered complete? *Physical Review* 47(10): 777–780.
- Fine, K. 2005. *Modality and Tense*. Oxford: Oxford University Press.
- Hofer-Szabó, G., M. Rédei, and L.E. Szabó. 2013. *The Principle of the Common Cause*. Cambridge: Cambridge University Press.
- Horty, J.F. 2001. *Agency and Deontic Logic*. Oxford: Oxford University Press.
- Horty, J.F., and N. Belnap. 1995. The deliberative stit: A study of action, omission, ability and obligation. *Journal of Philosophical Logic* 24: 583–644.
- Kamp, H. 1971. Formal properties of ‘now’. *Theoria (Lund)* 37: 227–273.

- Kaplan, D. 1989. Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals; and afterthoughts. In *Themes from Kaplan*, ed. J. Almog, J. Perry, and H. Wettstein, 481–563; 565–614. Oxford: Oxford University Press.
- Kracht, M., and O. Kutz. 2007. Logically possible worlds and counterpart semantics for modal logic. In *Philosophy of logic, Handbook of the philosophy of science*, ed. D. Jacquette, 943–995. Amsterdam: Elsevier.
- Lewis, D.K. 1986. *On the plurality of worlds*. Oxford: Blackwell.
- MacFarlane, J. 2003. Future contingents and relative truth. *The Philosophical Quarterly* 53(212): 321–336.
- MacFarlane, J. 2014. *Assessment sensitivity: Relative truth and its applications*. Oxford: Oxford University Press.
- Mackie, J.L. 1980. *The cement of the universe. A study of causation*. Oxford: Oxford University Press.
- Malpass, A., and J. Wawer. 2012. A future for the thin red line. *Synthese* 188(1): 117–142.
- McCall, S. 1990. Choice trees. In *Truth or consequences. Essays in honor of Nuel Belnap*. J. Dunn and A. Gupta, eds., 231–244. Dordrecht: Kluwer.
- McCall, S. 1994. *A model of the universe*. Oxford: Oxford University Press.
- Müller, T. 2002. Branching space-time, modal logic and the counterfactual conditional. In *Non-locality and Modality*, ed. T. Placek and J. Butterfield, 273–291. Dordrecht: Kluwer.
- 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. Branch dependence in the “consistent histories” approach to quantum mechanics. *Foundations of Physics* 37(2): 253–276.
- 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, general relativity, the Hausdorff property, and modal consistency. Technical report, Theoretical Philosophy Unit, Utrecht University. <http://philsci-archive.pitt.edu/8577/>.
- Müller, T. 2011b. Probabilities in branching structures. In *Explanation, prediction and confirmation. The Philosophy of Science in a European Perspective*, Vol. 2, eds. D. Dieks, W.J. Gonzalez, S. Hartmann, T. Uebel and M. Weber, 109–121. Dordrecht: Springer.
- Müller, T. 2013a. Alternatives to histories? Employing a local notion of modal consistency in branching theories. *Erkenntnis*. doi:10.1007/s10670-013-9453-4.
- Müller, T. 2013b. A generalized manifold topology for branching space-times. *Philosophy of science*, forthcoming; preprint URL = <http://www.jstor.org/stable/10.1086/673895>.
- Müller, T., N. Belnap, and K. Kishida. 2008. Funny business in branching space-times: infinite modal correlations. *Synthese* 164: 141–159.
- Penrose, R. 1979. Singularities and time-asymmetry. In *General relativity: an Einstein centenary survey*, ed. S.W. Hawking and W. Israel, 581–638. Cambridge: Cambridge University Press.
- Placek, T. (2010). Bell-type correlations in branching space-times. In *The Analytic Way. Proceedings of the 6th European Congress of Analytic Philosophy*, eds. T. Czarnecki, K. Kijania-Placek, O. Poller and J. Woleński, 105–144. London: College Publications.
- Placek, T., and N. Belnap. 2012. Indeterminism is a modal notion: branching spacetimes and Earman's pruning. *Synthese* 187(2): 441–469.
- Placek, T., and L. Wroński. 2009. On infinite EPR-like correlations. *Synthese* 167(1): 1–32.
- Ploug, T. and P. Øhrstrøm. 2012. Branching time, indeterminism and tense logic. Unveiling the Prior-Kripke letters. *Synthese* 188(3): 367–379.
- Pörn, I. 1977. *Action theory and social science: Some formal models*. Dordrecht: D. Reidel.
- Prior, A.N. 1957. *Time and modality*. Oxford: Oxford University Press.
- Prior, A.N. 1967. *Past, present and future*. Oxford: Oxford University Press.
- Quine, W. 1980. *From a logical point of view. Nine logico-philosophical essays*. Cambridge: Harvard University Press. (2nd, revised edition.)

- Rumberg, A., and T. Müller. 2013. Transitions towards a new account of future contingents. (Submitted.)
- Steward, H. 2012. *A Metaphysics for Freedom*. Oxford: Oxford University Press.
- Strobach, N. 2007. *Alternativen in der Raumzeit: Eine Studie zur philosophischen Anwendung multimodaler Aussagenlogiken*. Berlin: Logos.
- Szabó, L., and N. Belnap. 1996. Branching space-time analysis of the GHZ theorem. *Foundations of Physics* 26(8): 989–1002.
- Thomason, R.H. 1970. Indeterminist time and truth-value gaps. *Theoria* 36: 264–281.
- Thomason, R.H. 1984. Combinations of tense and modality. In *Handbook of philosophical logic, vol. II: extensions of classical logic*, volume 165 of *Synthese Library, Studies in Epistemology*, eds. D. Gabbay and G. Guenther, 135–165. Dordrecht: D. Reidel Publishing Company.
- Thomson, J.J. 1990. *The realm of rights*. Cambridge: Harvard University Press.
- Van Inwagen, P. 1983. *An essay on free will*. Oxford: Oxford University Press.
- von Kutschera, F. 1986. Bewirken. *Erkenntnis* 24(3): 253–281.
- von Wright, G.H. 1963. *Norm and action. A logical inquiry*. London: Routledge.
- Weiner, M., and N. Belnap. 2006. How causal probabilities might fit into our objectively indeterministic world. *Synthese* 149: 1–36.
- Wroński, L., and T. Placek. 2009. On Minkowskian branching structures. *Studies in History and Philosophy of Modern Physics* 40: 251–258.