


Facing the Future

Agents and Choices in
Our Indeterminist World



Nuel Belnap
Michael Perloff
Ming Xu

Facing the Future

This page intentionally left blank

FACING THE FUTURE

Agents and Choices in Our Indeterminist World

NUEL BELNAP

MICHAEL PERLOFF

MING XU

With Contributions by

Paul Bartha

Mitchell Green

John Horty

OXFORD
UNIVERSITY PRESS

2001

OXFORD

UNIVERSITY PRESS

Oxford New York
Athens Auckland Bangkok Bogotá Buenos Aires Cape Town
Chennai Dar es Salaam Delhi Florence Hong Kong Istanbul Karachi
Kolkata Kuala Lumpur Madrid Melbourne Mexico City Mumbai Nairobi
Paris São Paulo Shanghai Singapore Taipei Tokyo Toronto Warsaw
and associated companies in
Berlin Ibadan

Copyright © 2001 by Oxford University Press, Inc.

Published by Oxford University Press, Inc
198 Madison Avenue, New York, New York 10016

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
electronic, mechanical, photocopying, recording, or otherwise,
without the prior permission of Oxford University Press

Library of Congress Cataloging-in-Publication Data

Belnap, Noel D., 1930-

Facing the future : agents and choices in our indeterminist world /

Noel Belnap, Michael Perloff, Ming Xu, with contributions by Paul Bartha,
Mitchell Green, John Horty.

p cm

Includes bibliographical references and index.

ISBN 0-19-513878-3

1 Agent (Philosophy) 2. Choice (Psychology) 3 Free will and determinism

I Perloff, Michael II Xu, Ming III Title

B105 A35 2001

128' 4—dc21 00-064995

9 8 7 6 5 4 3 2 1

Printed in the United States of America
on acid-free paper

Agents and choices in branching time with instants

We continue to discuss a variety of foundational matters concerning the theory of agents making choices in our indeterminist world.* Other chapters, especially chapter 2, have set out and briefly discussed a number of postulates—the $BT+I+AC$ postulates—on which stit theory and its semantics is based. In this chapter we consider these postulates one by one. We begin in §7A by treating at some length the postulates concerning branching time with instants—the $BT+I$ postulates. The section includes a rudimentary discussion of the interplay between “propositions” and “events” in $BT+I$ theory, and considers alternative theories, especially $T \times W$ theory, that are relevant to indeterminism. We proceed in §7B to offer some reflections on indeterminism, to take up determinism and its denial, to consider arguments against branching, and to point out how “causality” arises naturally in branching time. §7C discusses those postulates of $BT+I+AC$ theory that relate to agents and their choices. In §7D we discuss the postulate introducing the domain of quantification.

7A Theory of branching time

C. P. Snow rightly says that our academic society falls apart into two cultures, that of the scientist and that of the humanist. There is nevertheless but one world, the common home of physical process and of agency. To say as much is not to become embroiled in reductionist or antireductionist disputes; it is simply to note quasi-geometrical facts such as that Jack’s very human deliberation about whether to go to the beach *preceded* the very physical destruction

*With the permission of Kluwer Academic Publishers, parts of this chapter draw on Belnap 1996a.

of the beach by the hurricane, which happened *far away from* (and in causal independence of) Alfredo's finishing his pasta primavera, which was followed by a coin flip that by chance came up heads. It is in this spirit that we propose branching time with agents and choices as a high-level, broadly empirical quasi-geometrical theory of our world that counts equally as proto-physical and proto-humanist. Like geometry, it does not pretend to be that famous "theory of everything." It concerns above all the structural aspects of how the doings of agents fit into the indeterministic causal structure of our world. We postulate an underlying temporal-modal-agent-choice structure $\langle \textit{Tree}, \leq, \textit{Instant}, \textit{Agent}, \textit{Choice}, \textit{Domain} \rangle$ subject to certain constraints. You will find a list of all "*BT+I+AC* postulates," with *BT+I+AC* standing for "agents and choices in branching time with instants," in §3. We go over these postulates one at a time, concentrating on foundational questions.

7A.1 Nontriviality and partial order

We take up the two most basic postulates before going on to those that are distinctive of branching time as a theory of our world. The first postulate is especially trivial.

NONTRIVIALITY. (*BT+I+AC* postulate. Reference: **Post. 1**) *Tree* is a nonempty set: $\textit{Tree} \neq \emptyset$.

Tree is intended to represent our world as a set of events, so that nontriviality says that something happens. Its role is technical; we do not care to argue for the thesis that there is something rather than nothing. We are explicit that this representation, as such, gives no information about other sorts of things that our world doubtless contains; in particular, it is irrelevant to the purposes of this book whether events, or entities of some other category, are "fundamental" or "derived," nor does this book tackle the problem of how events relate to, for example, enduring objects such as brains and persons.

MOMENTS. (*Definition.* Reference: **Def. 1**) A *moment* is defined as a member of *Tree*. We let m and w range over moments; and we let M range over sets of moments.

Each moment should be pictured as an instantaneous, spatially unlimited, really possible concrete event ("super event" in the language of Thomson 1977), taken pre-relativistically. We intend that each moment stretches across all of space-time. This makes no relativistic sense, and at least because it is not relativistic, branching time is not a true theory. It is, however, an approximation to the truth, and will give us some modest insight into how our world hangs together. A more adequate proto-physical theory called "branching space-time" is presented in Belnap 1992. That theory takes seriously that physical and human happenings are local events rather than universe-wide. We nevertheless remain here with branching time, since no study known to us has attacked the presumably more difficult problem of understanding agency in anything like branching space-time.

Tree, although a set of moments, may perhaps be taken to represent or supervene upon a concrete entirety or world in something like the sense of Lewis 1986. It may matter whether the fundamental notion represented by *Tree* is concrete or abstract or mixed, or indeed whether the distinction itself should be taken as intramundane, but not for present purposes. What does matter is the key difference between this concept of "world" and that of Kripke 1959 (etc.), of Lewis 1986 (etc.), or of the standard four-dimensional concept derived (we suppose) from Newton by way of Einstein and Minkowski: The world, our (only) world, contains real possibilities both for what might be and for what might have been. The fundamental idea is that possibility—real possibility, objective possibility—is in the world, not otherworldly. If there are alternative worlds, then they, too, come with their real possibilities.

Other philosophers take possibilities to be what is consistent either with the laws of logic, or with the laws of physics, thus giving possibilities a fundamentally linguistic status. Still others think of possibilities as abstract or as creatures of the mind. We do not mean to lodge an objection against any of the many other concepts of "possibility" suggested by philosophers in the course of other studies. We only urge that fashioning a rigorous theory of agency and indeterminism is worthwhile, and that in doing so it is greatly useful to construe possible events as both concrete and objective. This study presupposes, but does not argue for, this point of view.

Some thinkers doubt that it makes sense to think of a concrete event as necessarily (i) super-large and (ii) instantaneous. Of course none of our postulates guarantees this interpretation of "moment," but it is certainly what we intend. As for (i), although with Thomson 1977 we find super events entirely respectable, there are surely smallish events as well, to be studied in another theory such as the "branching space-time" theory mentioned earlier. We are careful not to claim insight into the special problems raised by considering local events. We take (ii) as either literally true or as a helpful idealization, following Euclid's use of "point." Either way, we have not the slightest objection to theories that work out a rigorous ontology of our world based on "intervals" or the like, provided they are rigorous and provided they are helpful. Evenhandedly, we do not count such theories as objections to ours. In a similar vein, it is critical to our enterprise, but not to other eminently worthwhile enterprises, that the theory advanced aspires to live up to the standards of rigor embodied in the work of Frege, with every concept sharply defined in terms of a family of primitives clearly stated to be such.

Sometimes we say "*Tree*" and sometimes we say "*Our World*." We use "*Our World*" when emphasizing that our theory is about our world, and we use "*Tree*" when we wish to call attention in a more abstract way to the particular quasi-geometrical properties of our world, its "shape," so to speak, that we soon begin postulating.

Tree is not just a set; it comes with an order on it that we call the "causal order."

CAUSAL ORDER. (*BT + I + AC postulate. Reference: Post. 2*) *Tree* is partially ordered by \leq :

Reflexivity. ($m \leq m$).

Transitivity. ($m_1 \leq m_2$ & $m_2 \leq m_3$) \rightarrow $m_1 \leq m_3$.

Antisymmetry. ($m_1 \leq m_2$ & $m_2 \leq m_1$) \rightarrow $m_1 = m_2$.

We immediately introduce $<$ as the irreflexive and asymmetric cousin of \leq .

STRICT CAUSAL ORDER. (*Definition. Reference: Def. 2*) $m_1 < m_2$ is the strict partial order of *Tree* associated with the partial order \leq ; that is, $m_1 < m_2$ iff ($m_1 \leq m_2$ & $m_1 \neq m_2$). We refer to either of these as the *causal ordering* of *Tree*.

Picture the direction of $m_1 \leq m_2$, which flows from past to future, as generally forward from m_1 to m_2 (see Figure 2.1).

CAUSAL ORDER READINGS. (*Definition. Reference: Def. 2*) We read \leq and its converse with the plain words *earlier/later, below/above, lower/upper, backward/forward*, and so on, and insert *proper* when we intend $<$ or its converse. When m_1 is properly earlier than m_2 , we also say, in a fashion much more revealing of our intentions, that m_1 is *in the (causal) past of* m_2 ; and m_2 is *in the (causal) future of possibilities of* m_1 .

Observe that this intuitive use of "causal" does not imply that the earlier moment *causes* the later one; the usage comes from those dealing with special relativity, a group that also does not pretend that *propter hoc* follows from a mere *post hoc* in the causal order. We need to use the language of causal order in part to enforce that we are not speaking of abstract "times," but rather of possible *concrete events*.

Reflexivity of \leq is, as always, a postulate of convenience; self-causation is not implied, and indeed is denied by the irreflexivity of $<$.

Transitivity in **Post. 2** is perhaps arrogant since it extends the causal order in a simple way to the furthest reaches of our world. That the postulate is something like the truth for the middle-sized portions of our world that lie close to hand, however, seems beyond doubt, so that unless one is engaged in studying either the very small or the very large, it seems a safe enough postulate. We certainly do not ourselves know how to fashion a simple theory that does without it.

Antisymmetry already signals that the domain of application of branching time must consist in nonrepeatables (fully concrete events), not abstract situations or "states" of either "systems" or "times" such as nearly everyone with some training in physics is likely to think of. To ask which abstract "states" can follow which is not at all the same as asking which concrete *events* can follow which, and does not give rise to the same theoretical constraints. States can come again: $s_1 < s_2$ and $s_2 < s_1$, and hence $s_1 < s_1$, is admissible for states. Concrete events such as moments, however, are not "repeatables" and cannot properly precede themselves.

The following easily defined concept underlies much of our work.

COMPARABILITY. (*Definition. Reference: Def. 3*) Moments m_1 and m_2 are comparable if \leq goes one way or the other: $(m_1 \leq m_2)$ or $(m_2 \leq m_1)$.

We do not postulate outright that there are incomparable moments, but that certainly must be the case to give point to this study, whose entire thrust is to make sense out of a world in which we face incomparable possible events in our future of possibilities—that is to say, distinct worldwide events neither of which lies in the past of the other.

Next we introduce the notions of chain and history.

CHAINS AND HISTORIES. (*Definition. Reference: Def. 3*)

- A *chain*, c , in *Tree* is a subset of *Tree* such that every pair of its members is comparable: $c \subseteq Tree \ \& \ \forall m_1 \forall m_2 [m_1, m_2 \in c \rightarrow m_1 \leq m_2 \text{ or } m_2 \leq m_1]$. We let c range over chains in *Tree*.
- A *history*, h , of *Tree* is a maximal chain in *Tree*: h is a history of *Tree* iff h is a chain in *Tree*, and no proper superset of h is itself a chain in *Tree*. We let h range over all histories.
- *History* is the set of all histories of *Tree*. We let H range over subsets of *History*.

In branching time, chains represent certain complex concrete events, some of which we will categorize in §7A.4. History is, however, *the* essential concept. A history represents a single possible course of events. A history has run on from time immemorial, and will presumably run on forever. Nothing can consistently be added to a history, neither ahead, nor behind, nor in the middle. Histories take to the ends of time “a way events can go,” and are in this sense exceedingly “long.” Histories also take definiteness to the limit (they decide all disjunctions), and are, in contrast to a thickish family of histories, maximally “skinny.” Note, however, that we do not take histories as a *primitive* idea; given *Tree* as a set of moments ordered by \leq , the concept of history is already there—unless one suspects set theory, or has philosophical qualms with taking a commonsense idea to the limit. These suspicions or qualms may or may not in the end be warranted, but to take such (perfectly legitimate) philosophical concerns as blocking the road to even beginning an inquiry such as this one seems likely to interfere with one’s understanding of agents and their doings.

Perhaps the idea of a history as a *set* of moments should be taken to supervene on an underlying notion of history as a concrete whole that has parts instead of members, but it doesn’t matter for present purposes. What does matter is that a history is not an entirety or world. Our world is chock-a-block with real possibilities, perhaps with chances, and certainly with actions, and therefore with choices among sets of incompatibilities, none of which find a home in a single history.

Chains and other sets of moments can be bounded either above or below; for perhaps pedantic completeness, we assemble the standard definitions.

BOUNDS. (*Definition. Reference: Def. 5*)

- A moment, m , is a [proper] lower bound of a set of moments, M , iff m is [properly] earlier than every member of M ; and similarly for [proper] upper bounds of chains.
- We let $m < M$ iff m is a proper lower bound of M , and similarly in other cases.
- We let $M_1 < M_2$ iff every member of M_1 is earlier than every member of M_2 , and similarly in other cases.
- Greatest lower bounds and least upper bounds of sets of moments are as usual in the theory of partial orders.

7A.2 Forward branching only

Now that we have introduced histories, let us consider the difference between their forward and backward branching. We do not, as we have said, postulate forward branching, but we certainly expect it. Looking backward, however, we postulate that there is no branching toward the past. After explaining the postulate, and giving some consideration to forward branching, we argue in favor of no backward branching.

NO BACKWARD BRANCHING. (*BT + I + AC postulate. Reference: Post. 3*) Incomparable moments in *Tree* never have a common upper bound; or contrapositively, if two moments have a common upper bound, then they are comparable: $(m_1 \leq m_3 \ \& \ m_2 \leq m_3) \rightarrow (m_1 \leq m_2 \ \text{or} \ m_2 \leq m_1)$.

This is the postulate that makes *Tree* look like a tree. In order to help see its significance, it is convenient to define ideas of "past" and "future" with a Dedekind-like abstractness that does not presuppose any special order-type on the histories (other than the fact that they are chains).

CUTS, PASTS, AND FUTURES. (*Definition. Reference: Def. 6*)

- A *historical cut* for a history, h , is a pair $\langle p, f \rangle$ of sets of moments such that (i) neither p nor f is empty, (ii) $p < f$, and (iii) $(p \cup f) = h$.
- p is the *past history* of f iff f is a *future history* of p iff $\langle p, f \rangle$ is a historical cut.¹
- M is the *causal past* of a future history, f , iff M is the set of all proper lower bounds of f .
- We often say just *past* because given no backward branching, a causal past is the same as a historical past.

¹We use "p" for past histories and "f" for future histories, but also for other purposes; we believe, however, that no ambiguities arise in this book.

- M is a *future of possibilities*, or a *causal future* of a past, p , iff M is the set of *all* proper upper bounds of p .

In branching time, a future history must be distinguished from a future of possibilities, so that in this book we never say merely "future."²

- We say that $\langle p, M \rangle$ is a *causal cut* iff p is a past, and M is the future of possibilities of p .
- All of these usages may straightforwardly be adapted to speak of pasts and futures of single moments.

The postulate of no backward branching tells us that each future history has exactly one past history: If $\langle p_1, f \rangle$ and $\langle p_2, f \rangle$ are historical cuts, then $p_1 = p_2$. Where f is a future history, this uniqueness entitles us to introduce the phrase "*the past history of f .*" Since no backward branching implies that the past history of f is the set of *all* moments preceding f , by strict analogy with special relativity, we may also call the past history of f the *causal past* of f , or just its (plain) *past*.

As a special case of no backward branching, *our* causal past is a history in which everything has a place in terms of earlier/later; our past is not an assemblage of incompatible possibilities. If two possible events (i.e., members of *Tree*) each lie in our past, then one of them was causally earlier, and lies in the past of the other, giving rise to an actual linear causal sequence, earlier-later-now. This is not merely a matter of temporal dating; it concerns causal linkage. This discussion does, however, falsely but conveniently presuppose that causal linkage is between entire spatial "slices" instead of between small local events. As previously noted, this defect is remedied in Belnap 1992, but not in this book.

Although a future history has but one past history, the converse is by no means true. If determinism is not permanently true, there will be past histories p to which more than one future history can be appended: One can have historical cuts $\langle p, f_1 \rangle$ and $\langle p, f_2 \rangle$ with $f_1 \neq f_2$. There is therefore no "rigid" sense to the expression "the future history of p ." We may, however, reasonably speak of a unique "future of possibilities." Each future of possibilities will look something like a tree, and will be a subtree of *Tree*.

When at a given moment Lee-Hamilton says "the past is stone, and stands forever fast," his use of the phrase "the past" safely refers to a past that is uniquely determined by the (idealized) context of utterance.³ Furthermore, this past is a portion of each history of which his utterance is a part: No matter

²In earlier publications we sometimes used "future" as a short form of "future of possibilities." That was an expository error; in delicate discussions of indeterminism, it is better to avoid this contraction because of its confusing conflict with ordinary English usage. One should be resolute in speaking of either "future histories," or "future of possibilities," never of just "futures."

³It is a confusion to infer from this that "it was true that Q " implies "it is settled that it was true that Q "; instead, as we note at the very end of this section, one has only that "it was settled true that Q " implies "it is settled that it was settled true that Q ."

the future history, that past stands forever fast. Given no backward branching, talk of so acting as to "influence" the past is therefore just talk.

When, however, at a certain moment Tennyson says "I dipt into the future, far as human eye could see, / Saw the Vision of the world, and all the wonder that would be," we must be careful, especially when we learn that the Vision involves a sea battle. There is no philosophical problem if we interpret the poet as denoting his future of possibilities, the tree that fans out from the moment of his utterance in intricate arborescent profusion—a proper source of wonder indeed. Nor is there a problem if the poet is using "the future" as "the future history," but with its "nonrigid" denotation relative to a poetically visionary history—as long as we do not forget this relativization and take proper care to disambiguate Tennyson's expression accordingly, just as we would in the case of an occurrence of "the river" in *The Lady of Shalott*. (This simply adapts to singular terms the sentential semantic insight of Prior-Thomason that we adopted and first explained in §2A.1, and that we make more explicit in §8E.) If, however, by the phrase "the future" the poet intends to denote a *unique future history of which his moment is a part*, then his intention cannot be carried out, for (unless determinism be true from that moment as far as human eye can see) there is no such unique future history, as from time to time we repeat.

We believe (if that is the right word) in forward branching, and in the impossibility of backward branching. Sometimes one hears a philosopher or a physicist maunder on about distinct pasts that coalesce in a present moment, and doubtless it is good that our conceptual limits be tested. We confess, however, that we ourselves cannot follow these fancies. That we face alternative future histories seems to us right; that we are faced away from alternative pasts seems to us wrong. That starting with the concrete event that occurred yesterday morning there were incompatible possible events each of which might have transpired seems to us right; that more than one of these incompatible possible streams of events might have finished up in this very concrete situation seems to us wrong. In common with antisymmetry, no backward branching makes sense only for objective, concrete events. First, we advance no theory at all about what is possible (not objectively but) "for all one knows." A given concrete situation could obviously have been preceded by any of various inconsistent predecessors, "for all one knows." It is precisely to preclude this epistemic or doxastic use of "possible" that we so tiresomely repeat that our present concern is with "objective" possibilities. Second, no backward branching fails to apply to "states" or other repeatable carriers of partial information. There is no doubt whatsoever that a present "state" may be accessible from either of two earlier incompatible states. There is no doubt about this because there are so very many senses to the word "state." Surely there are physical "systems" with a favored family of "states" that branch only forward and not backward, others that branch only backward and not forward, others that are doubly deterministic in terms of their favored "states," and still others with more exotic structural properties or with no interesting properties at all. After all, everything happens. None of this, however, is relevant to our postulate of no backward branching. To discuss any of it is to change the topic.




Figure 7.1: Backward branching?

EXAMPLE. (*Backward branching with states*) Count position and acceleration as adding up to a "state," and, invoking an example suggested by Bressan, slide a disc along the floor. Friction will bring it to a stop. Its future is then determined as continuing in the stopped state, at least for a while, whereas from its stopped state there is no inference to when in the past it was thrown. True, but irrelevant. No backward branching does not imply that this particular definition of "state" gives us information as to when the disc was thrown; it only implies that regardless of the poverty or richness of any concept of "state" that is brought into play, there is a fact of the matter admitting no real alternatives. The concrete event of the disc coming to a halt has in its past a unique concrete event of its being thrown—a fact that is no less true for being absent from physical theories cast in terms of systems and states.

One of the reasons we think no backward branching a necessity is that we do not know how to make sense of agency without it. We expose our puzzlement with two side-by-side diagrams. In Figure 7.1, the left-hand diagram represents our own theory of a moment of choice. Picture a castaway on a deserted island. Let us position ourselves at moment m_1 , and let us suppose that at 1:00 P.M. the castaway chose between lighting his signal fire or not. We may now say that ever after 1:00 P.M. it has been a settled matter which choice the castaway made: He chose to light the fire. Because of his choice, at 3:00 P.M. he was rescued. We can say that although before 1:00 P.M. no-fire was a real possibility for the castaway, that possibility was not realized. Of course we at m_1 (perhaps in England) might not know which choice the castaway made, and we might not know whether or not he was rescued; but given our position at m_1 , there is a settled fact of the matter.

In contrast, the right-hand diagram in Figure 7.1 represents backward branching. Let us position ourselves at moment m_3 in this right-hand diagram. We can still say that the following is now settled at 2:00 P.M., and also at 3:00 P.M.: Either it was settled that the castaway chose fire or it was settled that the castaway chose no-fire. Accordingly, no matter which history through m_3 you pick, at those times it was definitely settled whether or not the castaway was rescued at 3:00 P.M. This is in accord with our usual views about the consequences of choice: Choosing the signal fire settles the matter in favor of rescue, choosing no-fire settles the matter in favor of its opposite, no-rescue. That sounds fine. The trouble is that now at 5:00 P.M. it is no longer settled which choice the castaway made, and it is no longer settled whether or not he was rescued. There used to be a settled fact of these matters, but it has dissipated. There is no longer a settled fact as to whether he was rescued or not at 3:00 P.M. The diagram seems to show that the castaway had a real choice at 1:00 P.M., but the same diagram seems to say that the choice at 1:00 P.M., real as it is (or was), has no differential consequences for us at 5:00 P.M. The consequence of choosing fire at 1:00 P.M. was rescue at 3:00 P.M., and the consequence of choosing no-fire at 1:00 P.M. was no-rescue at 3:00 P.M. But whatever status "rescued at 3:00 P.M." has for us at 5:00 P.M., on this diagram "not rescued at 3:00 P.M." has quite the same status.

Since the differential consequences of the choice have disappeared by 5:00 P.M., it may be that we should say that the choice has also disappeared: That the castaway had a choice at 1:00 P.M. used to be true (at 2:00 P.M. and 3:00 P.M.), but it isn't true any longer (at 5:00 P.M.). We hope that such talk makes you as nervous as it does us. We do not mean to have exhibited a contradiction, or even a manifest absurdity. We intend only to expose our puzzle about making sense of agency if there is backward branching. Perhaps backward branching is all right; we just don't see how. We do not know how to develop a sensible theory of how choices fit into the world except with the understanding that once a choice between options is made, for ever after it is a definite and settled matter which choice was made (see (*) on p. 39).

Here is another example. According to the present theory, if you choose on Saturday to promise your neighbor to mow his lawn, then ever afterward it is a settled matter that you made the promise. You may forget, your neighbor may forget, everyone may forget, but the fact remains a fact. Even if you are released from your promise, it will be in virtue of later happenings. You cannot hope for release by its becoming true that you didn't make that choice on Saturday, or even its becoming "possible" that you didn't make it. You may reasonably hope that the promise will be voided in any one of a number of ways; but no matter how long you wait in hope, the settled fact that on Saturday you chose to promise to mow the lawn is not going to go away. If, however, we assume that histories can branch backward, then such a hope is (perhaps) reasonable. And if so, it can happen that on Monday it is a settled fact that on Saturday you made a promise, and then, on Tuesday, the "fact" that you promised can become "possibly false" (or "not definitely true"—we hardly know what words to use to describe this imaginary situation that in fact we are unable to imagine). Maybe

that makes sense; but we ourselves have absolutely no idea (in the absence of doublethink) of how to develop a theory of promises in such a setting.

One pressure for backward branching comes from certain interpretations of certain physical theories. According to these interpretations of these theories, the world looks just the same upside down. We of course have no special expertise in physics, and so turn to the quantum-field theorist Haag 1990.

I want to suggest here that—once one accepts indeterminism—there is no reason against including irreversibility as part of the fundamental laws of nature. ... It should be stressed that this picture does not touch CTP-invariance or detailed balancing. ... the term “time reversal” should be replaced by “motion reversal.” (p. 247)

... the use of probabilities in statistical mechanics and quantum theory is necessarily always forward directed since the past is factual and the future open. If irreversibility is introduced on a fundamental level as proposed then the coincidence of the different “arrows of time” (psychological, thermodynamic, cosmological ...) is immediate and in particular dissociated from any cosmological model. (p. 250; citations omitted)

Evidently even physicists disagree about these matters; it is, however, surely respectable to take Haag as our governing authority. If we do so, backward branching receives no aid and comfort from any physical theory.

The tense-logical principle that seems to us most closely associated with no backward branching is this: “If something was settled true, then it is (now) settled that it was true.” This principle—which is not to be confused with its invalid sound-alike, “If something was true then it is (now) settled that it was true”—is the one whose loss would leave us with a feeling of having lost our grip. Settled features of past moments stay settled. In particular, if we look back at earlier choice points, it is and will always remain a settled matter which choice was made at that earlier point.

7A.3 Historical connection

The $T \times W$ theories discussed in §7A.6 demand from the outset that no event belong to more than one history, that no event have more than one possible future in the past of which it lies. The next postulate not only permits meaningful overlap of histories, but insists on it.

HISTORICAL CONNECTION. (*BT + I + AC postulate. Reference: Post. 4*) Every two moments have a lower bound: $\forall m_1 \forall m_2 \exists m_0 [m_0 \leq m_1 \ \& \ m_0 \leq m_2]$. In other words, every two histories intersect.

So every two histories share a common past: their nonempty set-theoretical intersection. Our world, as represented by *Tree*, has no loose or floating pieces; it constitutes a whole, bound together by the causal order. It is perhaps this

postulate, above all, by means of which the theory insists that we not admit histories that are idle creatures of the imagination, or histories that could be "defined" by piecing together some arbitrary array of logical or conceptual or scientific possibilities. Instead, every possible history, h , has a definite causal relation to the very moment in which we converse, since h must share with it a common past. Historical connection is the postulate that endows the theory with a sense of robust reality. In other words, historical connection makes *Tree* a single world, *Our World*, instead of a mere collection of "worlds." In this sense it is historical connection that ensures that agents are actual and that our choices are real.

Historical connection puts all moments of *Tree* "in suitable external relations," as Lewis 1986 says (p. 208). It gives content to exactly the sort of real possibility that is pertinent to an understanding of stit. Just to make things clear by an example, we are disallowing that it is or was really possible that there should be blue swans unless there is some definite moment in our past that has a moment in its future of possibilities at which there are blue swans. Naturally, as armchair philosophers of indeterminism, we do not claim special insight into what is really possible; that is a matter for common sense or science or metaphysics. The point of the example is only to express our doubt that it is *easy* to be sure that it is or was really possible that there should be blue swans. Of course something terminological is going on here: We are using "really possible" as what is or was determined as possible in the world of our context of utterance, and thus in a sense much narrower than that sought by, for example, Lewis 1986 through the idea of recombination. But there is also something nonterminological: We think that the Humean picture of enormous recombinational possibilities for the immediate future (e.g., blue swans on our desk one nanosecond from now—Lewis 1986, p. 91, says that "anything can follow anything") is not relevant to what can be seen to, and that instead what counts is only the current—much narrower—set of real possibilities.

When the intersection $h_1 \cap h_2$ of two (distinct) histories h_1 and h_2 is not only nonempty, as promised by historical connection, but has a least upper bound, m_0 , we say that h_1 and h_2 *split* at m_0 (**Def. 7**). We write $h_1 \perp_{m_0} h_2$ (**Def. 4**). And if every two (distinct) histories split at some moment, we say that the *semi-lattice condition* is satisfied (**Def. 7**), since in context this holds iff every two moments have a greatest lower bound. Is it true? Do each two two histories split at a moment? In the branching space-time theory of Belnap 1992 we postulate the principle of "prior choice," which says that if a point-event, e_1 , belongs to one history, h_1 , and not to another history, h_2 , then there is a particular point-event, e_0 , in the past of e_1 such that h_1 and h_2 split at that point. In the context of branching space-time prior choice is, as far as we can see, a deep causal principle, and in fact in that context, we do not know how to carry on without it. The same words in the context of branching time express the semi-lattice condition. We are therefore certainly tempted to enter that postulate as a strengthening of historical connection, especially because there is no reason of which we know to suppose that the semi-lattice condition is false. On the other hand, by careful formulations that do not rely on that

condition, one can come to see that it does not make any difference to the theory of agents making choices in branching time. The fundamental reason for this seems to be that in our theory of choice, we postulate that regardless of the semi-lattice condition, when an agent makes a choice, there is always a last *moment* of indetermination, which is therefore the greatest lower bound of two moments belonging to different choices. In other words, if the fact that we are in one history rather than another is to be explained by *choice*, then those two histories, anyhow, split at a *moment*.

All this abstruse talk is to explain why we waver when it comes to postulating the semi-lattice condition. In the end we decide somewhat arbitrarily to do without it. Wavering to one side, however, even in the rudimentary context of branching time, one ought to see historical connection, especially when strengthened to the semi-lattice condition, as a powerful causal principle that, like prior choice, says that a real cause for something being one way rather than another always lies in the past. If histories were disconnected, one could not say that without bringing in soft ideas such as "similarity."

We also refrain from postulating that all the moments in *Tree* have a common lower bound, which would imply a kind of nonrelativistic Big Bang.

7A.4 Propositions, events, and their interplay

Branching time makes possible several important functional connections between concrete events on the one hand and "propositions" on the other, connections that seem difficult to make clear in other frameworks. This material is seldom used explicitly in this book, and may be skipped. Nevertheless, the ideas are implicit in many of our formulations, and explicit in a few, so that we take a bit of space to lay them out.

Propositions. Let us begin by articulating the idea of "proposition." On our view, "proposition" is serviceable but inconstant philosophical jargon. That is, in spite of its usefulness, there should never be an assumption that we all use it in the same way, or that anyone knows what anyone else means. "What is 'grasped' by a mind" doesn't help to pin down what a philosopher might mean, since "grasp" is an unpacked metaphor. "Meaning of a sentence" helps only to the extent that a particular understanding of "meaning" is in play. We stick with some "intensional" notion of propositions closely tied to the idea of (parameterized) "truth-conditions," thereby giving up any pretense of using an "intentional" concept that is essentially suitable for mental "grasping."

Having limited our target to propositions in a sense that resonates with parameterized truth conditions, we next observe that among those who think about time, there is sometimes a debate concerning whether a proposition can "change its truth value." We are sure that among the various concepts of proposition, some make such a change sensible, and some do not.

- If the proposition is taken to be a "timeless" or "unchanging" proposition, one may represent it in branching time by means of a set of histories,

H , thinking of H as *true in* just its members. This already makes the conceptual role of histories in the theory of branching time in some deep respects analogous to the role of "worlds" in Kripke-inspired modal logics. This notion of "proposition" is best tied to sentences considered as stand-alone and as assertable, (§6B.4.2), since such sentences are open only in the history of evaluation (combine Semantic thesis 6-5 and Assertability thesis 6-7).

- If one wishes to think of a proposition, however, as changing its truth value over time, one represents the proposition well in branching time by a set of moment-history pairs, that is, a subset of *Moment-History*. This "time-dependent" or "moment-dependent" notion of "proposition" is suitable for thinking about sentences considered as embedded in, for example, tense connectives.
- Indeed, when thinking of the contribution of sentences embedded in quantifier connectives, it can be helpful to advert to "assignment-dependent" propositions.

All these representations and concepts of propositions are useful. As a practical matter, and with no intent to legislate, we will reserve the unmodified notion of "proposition" for the first case: a set of histories, H . A proposition, H , is said to be *true* in each of its members, and *false* in each of its nonmembers. By itself this is trivial; we shall see, however, that the usage coheres with a number of other ideas. First we say a few words about events.

Events. Branching time permits only a rudimentary theory of events. That is because its primitive notion of a "moment" is already intended as something that is, although instantaneous, unbounded in spatial extent. Therefore no basis is provided for *local* events; for such events one would need to turn to branching *space-time* (Belnap 1992). It is good to think of a moment as a possible event, a possible *momentary* event. Momentary events automatically have their locus in the causal structure of our world, so that it makes sense to think of them as concrete.

A more general notion of a (concrete possible) event is as follows: A *concrete possible event* is a nonempty set, M , of moments. This Quine-like construction yields a kind of *de re* concept, since it gives only the causal locus of the event in our world. On the other hand, by antisymmetry, moments resemble Leibnizian monads insofar as each moment determines its entire past and its entire future of possibilities, a property that in some sense is passed on to sets of moments. In any case, we shall see that the theory of events qua sets of moments is more fecund and explanatory in branching time than it would be in an objective determinist theory.

What we say about a set of moments qua event depends on its causal "shape." Two cases are of special interest: "initial" events and "outcome" events. (Even when idealization is maximized, we do not propose that every set of moments

corresponds to an event in some intuitive sense. It is always the special cases that are of interest.) Consider the *setting up* or *initializing* of a situation that is indeterministic either by choice or by chance. One might have in mind, for instance, the process leading up to the flip of a coin. (A flip is of course a spatially local event, a feature that we must ignore.) There must exist some moments before it is decided whether heads or tails will eventuate, and also some moments after. Extend a series of pre-decision moments forward until further extension would have to involve a moment at which either heads is decided, or tails is decided. Let I be a set of moments resulting from this procedure; a set of moments such as I is what we call an "initial" event. (We use " I " as a mnemonic variable ranging over initial events.)⁴ Speaking abstractly, because, for example, heads-or-tails is by assumption eventually decided, the most essential condition on the causal "shape" of an initial is that it must be nonempty and properly upper bounded (by some moment at which the matter is decided), and therefore I must be a nonempty chain. For many purposes neither the "beginning" of such a setting-up nor its internal structure (e.g., denseness) makes a difference and may be ignored. For these purposes, we may therefore enter the following.

INITIAL EVENTS. (*Definition. Reference: Def. 8*) I is an *initial event* iff I is a nonempty and upper-bounded chain.

Fixing on an initial, I , as just defined, consider now a particular *outcome* of some indeterministic set-up, say "the coin's having come up Heads." Being an outcome, such an event, O , must be nonempty and causally lower bounded; so much is essential, for otherwise O could not come to be.

OUTCOME EVENTS. (*Definition. Reference: Def. 8*) O is an *outcome event* iff O is a nonempty and lower-bounded chain.

There are likely a number of enlightening notions of "outcome" obtainable from this beginning. Because the whole matter is, however, somewhat tangential to agency, we do not elaborate; we define only the following relational idea of one event being an outcome of another.

⁴Suppose the semi-lattice condition. Then an initial event could be defined more simply as a single moment, the moment at which some tails history splits from some heads history. Such a moment would be the last moment of indetermination. If it seems farfetched that there should be such a moment, observe that when, e.g., a ball begins to move, elementary physics infers from continuity principles that there is a last moment of rest, but no first moment of motion. To give up this bit of physics is to give up idealizing motion and therefore to give up on rigorous physical theory; to give up the existence of a last moment of indeterminateness (given the semi-lattice condition) is analogous.

A further analogy is also helpful. If we take seriously that the ball has parts that are space-like related in the sense of special relativity, then we cannot find a frame-invariant last moment of rest. A typical response to this problem is to idealize the ball as point-like, that is, as without space-like related parts. One does not need to believe that the ball "really is" point-like in order to profit from the clarity that such an idealization permits. We intend our idealized idea of a point-like upper end of an initial event to be just like that: You do not have to believe that initial events "really are" like that in order to profit from our proposed idealization.

OUTCOMES OF EVENTS. (*Definition. Reference: Def. 8*)

- O is a (possible) outcome of I , iff I is an initial event and O is an outcome event, and (every member of) I properly precedes (every member of) O .
- O is an immediate outcome of I iff O is a (possible) outcome of I , and if furthermore no moment lies properly between I and O .

The idea of a “transition” as an ordered pair of events now falls into place.

TRANSITIONS. (*Definition. Reference: Def. 8*) $\langle I, O \rangle$ is an [immediate] transition iff O is an [immediate] outcome of I .

The suggestion that it is sometimes better to construe events qua “happenings” as transitions rather than as events in our technical set-of-moments sense is made by von Wright 1981. Here we see that we are forced in this direction by the example of immediate transitions, where there is a “happening” but no room for an “event” (in the set-of-moments sense) between initial and outcome. In kinematics, the transition from rest to motion is a “happening” of this kind: transition, not event. The same is true for our idealized concept of choice considered as a happening. (Belnap 1999 expands on this thesis.)

Propositions from events. Having defined propositions as sets of histories and events as sets of moments, we briefly indicate some interrelations by means of definitions intended to be revelatory.

The proposition that such-and-such an event “exists” or “occurs” is much thrown around, usually in a contingent sense. Either in determinism or in Lewis-like constructions, however, there appears to be no objective, rigorous theory on the basis of which to make sense of these words. With unavoidable artificiality, we will say that an event “exists” in a contingent but timeless sense that relates the existence of an event to histories. (By implication, we reserve “occurs” for a moment-dependent idea; but we do not cash in this reservation.) It turns out that the sense to be given to “exists” rightly depends on the “causal shape” of the event that is presupposed. We define three notations that encode three different ways of mapping events into timeless “existence” propositions.

FROM EVENTS TO EXISTENCE-PROPOSITIONS. (*Definition. Reference: Def. 4*)

- $H_{(m)}$ is the set of histories in which m lies (or the set of histories “passing through” m): $h \in H_{(m)}$ iff $m \in h$.
- $H_{[M]} = \{h: M \subseteq h\}$, so that $H_{[M]}$ is the set of histories entirely containing M . For suitable M , this is the set of histories in which M passes away.
- $H_{\langle M \rangle} = \{h: (M \cap h) \neq \emptyset\}$, so that $H_{\langle M \rangle}$ is the set of histories that pass through at least one member of the set of moments, M . For suitable M , this is the set of histories in which M comes to be.

$H_{(m)}$ represents the contingent proposition that m timelessly *exists*, a proposition that is true in all and only those histories that m inhabits.

$H_{[M]}$ does not always make useful sense as a proposition. When, however, I is an *initial* event, $H_{[M]}$ is well interpreted as the (timeless) proposition that I is "completed," "finishes," or in Aristotle's phrase, "passes away." If you are thinking about experimental preparations or agonizing deliberations, *this* is a sensible thing to mean by "exists": A preparation-event or deliberation-event does not "exist" in a history unless it does so entirely. If some history, h , splits off in the middle of a deliberation, we decline to say that the deliberative event "exists" in h . (One would need explicitly to tackle the present progressive in order to have a rigorous account of a deliberation-event "existing" in histories in which it is not completed.) Note a comfortable interaction: The proposition that I is completed depends not at all on the causal "shape" of I in its nether region. Any initial cofinal with I toward the future will determine the same set of containing histories as does I . For example, if I is taken as a representation of the setting-up of an indeterministic experimental situation, whether or not the set-up is completed is insensitive to everything except the causal locus of its being completed.

When O is an *outcome* event, the meaning of "exists" needs to be quite different. In that case, the proposition $H_{(O)}$ is the right sense of "exists," since outcomes exist by beginning (not ending), and $H_{(O)}$ is true in all and only those histories in which O begins, commences, or in Aristotle's language, comes to be. Suppose a coin is flipped. If we look only at the bare causal structure, it is natural (but not necessary) to think of the flipping as lasting as long as there is indetermination as to the outcomes heads or tails, no matter how that flipping may be related to the way that the coin dances in the air. In the same way, it is natural (but not necessary) to think of the heads outcome beginning whenever the possibility of tails is excluded. In other words, the heads outcome (not must but) may be taken to begin whenever that outcome is determinately settled, and the proposition that the heads outcome exists should be true in just those histories in which settled-heads begins to be. For purposes of causal analysis, it doesn't matter when the heads outcome ends.

Consistency. Without belaboring the point, we note that the interplay between propositions and events in branching time generates a firmly based family of consistency/inconsistency concepts. We may take from possible-worlds theory the idea that when propositions are represented as sets of histories, *consistency* between them is definable as having some history in common, a history in which both propositions are true. We *add* that it is helpful to define the consistency of two initial events I_1 and I_2 by the consistency/inconsistency of the *propositions* $H_{[I_1]}$ and $H_{[I_2]}$, so that the question is whether or not there is a history in which both preparations are (not just started but) *completed*. Dually, two outcome events O_1 and O_2 are consistent/inconsistent iff the propositions $H_{(O_1)}$ and $H_{(O_2)}$ are consistent/inconsistent, so that the question is whether or not there is a history in which both outcomes begin to be. Even more delicately,

one may ask whether a certain initial is consistent with a certain outcome by way of well-chosen propositions. Finally, one has the right definition of what makes a transition “contingent.”

CONTINGENT TRANSITION. (*Definition. Reference: Def. 8*) $\langle I, O \rangle$ is a *contingent transition* iff $\langle I, O \rangle$ is a transition, and if some history is dropped in passing from the completion of I to the beginning of O : $H_{[I]} - H_{\langle O \rangle} \neq \emptyset$.

These remarks are intended to indicate with almost excessive brevity that branching time permits and indeed suggests a rigorous and modestly enlightening theory of the interrelations of (possible) concrete events and propositions. We mention that some additional ideas on initials and outcomes are offered in Szabo and Belnap 1996 and in Belnap 1996c, and in the unpublished set of notes Belnap 1995.

7A.5 Theory of instants

If “histories” are a way of making a sort of vertical division of *Tree*, then *Instant*, whose members are instants, is a kind of horizontal counterpart. In branching time, the doctrine of instants harkens back to the Newtonian doctrine of absolute time—and therefore is suspect. We use it, but we don’t trust it. For that reason if for no other we try to be as clear about it as we can. Instants are perhaps not fully “times” because we are not in this study relying on measures or distances, but it is intuitively correct to think of $i_{(m)}$ as the set of alternative possibilities for “filling” the time of m . We need instants because we think that for the achievement sense of stit, in considering whether Autumn Jane stit she was muddy at a certain moment, it is relevant to consider what else might have been at the instant inhabited by that moment. Evidently our uses of “moment” (in which we follow Thomason) and “instant” are jargon not sanctioned in ordinary speech, although the distinction is certainly there to be drawn.

Not all parts of stit theory rely on the theory of instants. Only the semantics for the achievement stit has need of these horizontal comparisons. For example, the theory of dstit, §8G.1, and the theory of strategies of chapter 13 are developed quite apart from the idea carried by *Instant*. In this important sense, the theory of instants is not a deep presupposition of stit theory. There is a contrast at this point with $T \times W$ theories as described in §7A.6. We nevertheless develop the theory of instants to the extent required by the semantics of the achievement stit. There are three postulates.

Instant AND INSTANTS. (*BT + I + AC postulate. Reference: Post. 5*)

- i. *Partition.* *Instant* is a partition of *Tree* into equivalence classes; that is, *Instant* is a set of nonempty sets of moments such that each moment in *Tree* belongs to exactly one member of *Instant*.
- ii. *Unique intersection.* Each instant intersects each history in a unique moment; that is, for each instant i and history h , $i \cap h$ has exactly one member.

w. Order preservation. Instants never distort historical order: Given two instants i_1 and i_2 and two histories h and h' , if the moment at which i_1 intersects h precedes, or is the same as, or comes after the moment at which i_2 intersects h , then the same relation holds between the moment at which i_1 intersects h' and the moment at which i_2 intersects h' .

We next offer some convenient definitions and simple facts, after which we comment on the postulates.

INSTANTS. (*Definition. Reference: Def. 9*)

- The members of *Instant* are called *instants*. i ranges over instants.
- $i_{(m)}$ is the uniquely determined instant to which the moment m belongs, the instant at which m "occurs."
- $m_{(i, h)}$ is the moment in which instant i cuts across (intersects with) history h ; that is, $i \cap h = \{m_{(i, h)}\}$.
- Order preservation can conveniently be stated in the symbols just introduced: $m_{(i_1, h_1)} < m_{(i_2, h_1)}$ implies $m_{(i_1, h_2)} < m_{(i_2, h_2)}$.
- Fact: $m_{(i_{(m_0)}, h_0)}$, a function of m_0 and h_0 , is the moment on history h_0 that occurs at the same instant as does m_0 : $i_{(m_{(i_{(m_0)}, h_0)})} = i_{(m_0)}$.
- $i|_{>m} = \{m_0: m < m_0 \ \& \ m_0 \in i\}$. We say that $i|_{>m}$ is the horizon from moment m at instant i .
- Where i_1 and i_2 are instants, we may induce a linear time order (not a causal order!) by defining $i_1 \leq i_2$ iff $m_1 \leq m_2$ for some moment m_1 in i_1 and some moment m_2 in i_2 . Instants can also be temporally (not causally) compared with moments, m : $i_1 < m$ iff $m_1 < m$ for some moment m_1 in i_1 ; and $m < i_2$ iff $m < m_2$ for some moment m_2 in i_2 .

Post. 5(i) encodes that making same-time comparisons between histories is objectively sound. We do not pretend to understand the conceptual problems involved in making such comparisons. The problem becomes ironically clearer when it is made more difficult by transference to branching space-time, where it is same-place-time comparisons between inconsistent point events that is at issue. All we can add is a conviction that it will not be possible to make suitable advances without consideration of the work of Bressan 1972, Bressan 1974, Zampieri 1982, and Zampieri 1982–1983, for they are the only persons we know who have worked within what seems to us the only reasonable position, that identifying place-times across possible situations is neither trivially easy (perhaps Kripke thinks this) nor a matter of partial constraint and partial stipulation (perhaps Lewis thinks this) nor empirically insignificant (perhaps this is van Fraassen's view), but a matter of serious physics.

Postulates **Post. 5(ii)** and **Post. 5(iii)** on *Instant* are very likely too strong (too oversimplifying); our justification is that agency is already hard to understand, so that it won't hurt to try to see what it comes to in circumstances that are not altogether realistic—as long as we keep track of what we are doing so that later we can try to move closer to reality. Thus, which most reality-oriented persons think not so plausible, but which greatly simplifies our picture of time, all histories are said to have exactly the same temporal structure. It follows that all histories are isomorphic with each other, and with *Instant*, which justifies the ordering on instants defined in **Def. 9**. On the other hand, no assumption whatsoever is made about the order type that all histories share with each other and with *Instant*. For this reason the present theory of agency is immediately applicable regardless of whether we picture succession as discrete, dense, continuous, well-ordered, some mixture of these, or whatever; and regardless of whether histories are finite or infinite in one direction or the other.

The theory of *Instant* is not, as we have said, as fundamental as that of \leq , and perhaps it is too strong, even pre-relativistically. Certainly the present assumption that all histories have isomorphic temporal orderings is stronger than comparable assumptions of Thomason 1970 or Thomason and Gupta 1980, and probably it should ultimately be weakened. In the meantime, while it is good to be concerned about oversimplification, the justification of our procedure is that the assumption can be clarifying when it comes to thinking about certain aspects of agency. *Instant* gives us a theory of *linear* “time” (based on \leq between instants) to play off against the theory of *branching* “time” (based on \leq between moments).

7A.6 Times \times Worlds and other alternative theories

Some of branching time can be modeled equally well on what might be called a “ $T \times W$ ” theory (with T for times and W for worlds) after Thomason 1984, or might be called a “divergence” theory after Lewis 1986. (See Zanardo 1996 and Di Maio and Zanardo 1998 for authoritative explanations and developments; this entire section has benefited from Zanardo's advice.) The key idea is that histories (to use the present terminology) are taken as ontologically distinct each from each (no common parts), and that enough in the way of additional concepts and argument is added to render it credible that two histories can “perfectly match up through” (Lewis) or “differ only in what is future to” (Thomason) a particular time. (Thomason in the passage cited is describing, not endorsing.) The modeling, however, is specious. The same bundle of (Humean) concepts and arguments that makes it credible that two histories might match *before* a certain time makes it equally credible that they match *after* a certain time while failing to match before that time; there is no asymmetry in the nature of things when it comes to mere matching. For this reason, such an approach to understanding our world, unlike branching time, can well encourage backward branching babbling. Branching time, incidentally, does not deny forward *matching* of distinct histories; instead what it denies is their forward *overlap*. For example, it may be that one of the histories on which we turn left becomes

very like one of the histories on which we turn right. Perhaps this happens at some point after the heat death of the sun; why not? But "very like" is not the same as sharing particular moments (and is not even very like). The gap between identity/distinctness on the one hand, and similarity/difference on the other, is unbridgeable. A helpful analogy is to points of space, some of which are well known as both very like and very distinct. The analogy is particularly apt since the theory of branching time is, in its general spirit, geometrical.

Every theory has merits and demerits. All $T \times W$ theories share the following two demerits: They need for their foundation a prior story about (i) times and (ii) worlds. Both of these stories are likely to be tall.

(i) The notion of a time as an entity independent of events gives pause. It would seem that in order to make sense out of indicating a particular time, one would have to explain quite a lot about clocks or cesium atoms or seasonal social practices, or something. This is something that $T \times W$ theorists need to do. The problem is perhaps best expressed in worrying over the requirement that one can make sense of "same time" across different worlds, which is normally entered as an unexamined and unquestionable presupposition of $T \times W$ theories. Of course in the theory of instants we rely on a postulate that is similar to this presupposition. That we introduce this postulate explicitly and separately after giving the fundamental ideas of branching time is conceptually significant, since we can develop the main ideas of tense and modality without it. The $T \times W$ approach hides this independence of tense and modality from the "same time" problem. Further, our "same time" postulation is explicit rather than presuppositional, so that it is easy to give it up when it does not help, or when it does not ring true.

In short, a $T \times W$ theory obviously provides no way to avoid presupposing instants or times shared across histories. Therefore the extent to which the concept of agency does not presuppose same-time comparisons across histories is the extent to which we have a good reason for not using $T \times W$ as a foundation for the theory of agency.

(ii) The idea of a possible world is the idea of something Very Big and (we suppose) hard to understand. As possibilities go, there are (we suppose) none bigger than worlds. It seems somehow a pity to start with something so big when what we want to understand is Caesar's situation when deciding whether or not to commence crossing the Rubicon. It would seem better to begin with a theory about more local incompatible possibilities, such as those available within ten or fifteen minutes, or available within ten or fifteen seconds, or (best) available immediately. The argument is that it is easier to credit a foundation built upon small possibilities, such as "the alternative moves in an actual chess game that I did not finally make" (Marcus 1985/86), rather than a foundation built upon something as large as worlds.⁵

Different from $T \times W$ is the framework employed in Chellas 1969, which one could call an " $h: T \rightarrow S$ " theory: Each history, h , is taken as a mapping from

⁵Marcus herself allows the point as persuasive, but argues against taking it seriously. It would be out of place to respond to her argument here.

the set T of times to a set S of states of affairs. An $h:T \rightarrow S$ theory allows that histories can “overlap” in the sense of sharing exactly the same states of affairs at exactly the same times, and perhaps, though we are unsure, this is genuine overlap. Still, on $h:T \rightarrow S$ theories, states of affairs are repeatables that can in principle occur at different times; they are therefore quite different from moments, each of which belongs to but a single instant, as is appropriate to a concrete event. Perhaps one could obtain a useful $h:T \rightarrow S$ theory from branching time by locating an interesting sense of “state of affairs” as partially characterizing moments. Otherwise the $h:T \rightarrow S$ theory seems faced with the double demerit of needing to provide both an account of times, and an account of states of affairs, on pain of being without application. Perhaps, for instance, belief in “states of affairs” is wild-eyed metaphysics (we hope not).

Thomason 1984 discusses the notion of a “Kamp structure” which is a variation on the $T \times W$ idea. (The literature uses the word “frame” more often than “structure”; there is no difference in concept.) In a Kamp structure, each world is provided with its own temporal ordering. No real ground is gained, however, since using these structures to simulate branching still requires the concept expressed by saying that two worlds “perfectly match up through” a given time. Perhaps ground is even lost, since the ontology of Kamp structures requires making sense of the possibility that the times of two worlds could be dramatically different in their ordering, while nevertheless sharing some particular entity ontologically classified as a “time.” For example, the respective times of two worlds might each contain both 4:00 P.M. and 5:00 P.M. The first world, however, could put 4:00 P.M. and 5:00 P.M. in their natural order, while the second world reverses their order. In considering what this could mean, we seem to lose our grip. In any case, all of our philosophical objections to $T \times W$ as a theory of the structure of our world apply equally to Kamp structures. (This discussion is indebted to correspondence with A. Zanardo.)

Both $T \times W$ structures and Kamp structures can be defined with first-order conditions. Yet another such representation of branching times is the notion of an “Ockhamist structure” $\langle X, <, \sim \rangle$ of Zanardo 1996. The idea is that one can represent the individual histories of *Tree* by letting $<$ be the union of disjoint linear orders, and one can represent the moments of *Tree* by making \sim be an equivalence relation standing for “same moment as.” We learn from Zanardo 1996 that Kamp structures, and Ockhamist structures, while exhibiting different ontologies, are mathematically equivalent representations. We shall complain about them equally, and about $T \times W$ as well. The complaint is, however, easiest to understand in terms of “bundled trees” and “moment-history structures,” each mathematically equivalent to Kamp and Ockhamist structures.

“Bundled trees” are closer to branching time in their ontology. (The idea is from Burgess 1978, 1979, 1980, the phrase from Zanardo 1996.) Bundled trees take “branches” as fundamental, where a *branch* is defined, using our terminology, as a forward-maximal chain (a future history) having a moment as its initial. This change in ontological perspective, while useful, is not, we think, of great importance. The reason is that branches are evidently in one-to-one correspondence with our moment-history pairs, and do approximately the

same work. There is, however, something about bundled trees that is indeed of considerable significance: Apart from perspective, bundled trees result from our trees $\langle Tree, \leq \rangle$ in exactly the same way that “Henkin general models” of second order logic result from standard models of second order logic (Zanardo 1996). Like second order logic, general models take the domain of individuals as primitive. Quite unlike second order logic, however, general models *also* take the range of predicate variables as primitive, which permits them to license the *omission* of some of the subsets of the domain of individuals, provided certain constraints are satisfied. Analogously, bundled trees take the set of branches to be primitive. This permits bundled trees to license the omission of branches, provided constraints are satisfied. When one adjusts for ontological perspective, this is tantamount to yet a fifth representation: One takes both the set *Tree* of all moments *and* the set *History* of all histories as primitive, in what might be called a *moment-history structure* $\langle Tree, \leq, History \rangle$. One is then in a position to license the omission from *History* of some of the histories (maximal chains) on $\langle Tree, \leq \rangle$, subject to the requirement that each moment belongs to at least one history. In this way one is able to replace second-order conditions with first-order conditions. Since first-order theories are always more technically tractable than second order theories, professional logicians tend to prefer to study them. This is by no means a misplaced “preference,” since the study of these more general moment-history structures provides a great deal of illumination in the way that, for example, much illumination is to be had by studying nonstandard first-order models of arithmetic, or of second order logic.

As a reading of Thomason 1984, pp. 151–152, suggests, however, there remains a question about the *descriptive adequacy* of the more general moment-history structures when taken as theories intended to be true to the facts of our world (or language?). In order to evoke the negative thrust of our opinion, we will label these more general structures as “missing-history structures.” An example will provide the reason that we think missing-history structures are descriptively inadequate. (After articulating the example we will make some remarks that examine how our negative opinion looks in the context of $T \times W$, leaving it to the reader to draw conclusions about Kamp structures, Ockhamist structures, and bundled trees.)

Let there be a specially interesting radium atom, a , such that as the seconds tick by after moment m_0 , the situation is as follows.

As long as a has not yet decayed, (i) a *might* decay before the next tick, and (ii) a *might not* decay before the next tick. (1)

We don’t need metrics in order to describe the situation, but it is essential to the story that the sequence of ticks has no upper bound. The situation is pictured in Figure 7.2, a diagram that we borrow from Thomason 1984. With reference to Figure 7.2, we let

$p \leftrightarrow$ atom a has not (yet) decayed,

so that




Figure 7.2: Suppose h_ω is not a history

$\sim p \leftrightarrow$ atom a has (already) decayed.

(Note that p is true or false independent of history, so that it makes sense to label the moments themselves with p or $\sim p$.) At issue is whether this situation can be properly described if the rightmost sequence, h_ω , is not considered to be a history. The missing-history approach says “yes,” whereas our no-missing-history approach says “no.”

The following is clear from Figure 7.2, as well as from (1):

Without exception, every no-decay chain (from m_0) of length n can be extended to a no-decay chain of length $n+1$. (2)

The “can” of (2) is not just mathematical. This “can” is to be taken in its usual historical-modal sense involving quantification over histories. In other words: At m_0 it is a settled fact that if a no-decay chain of length n will come to pass, then it is possible (but not guaranteed) that a no-decay chain of length $n+1$ will come to pass.

The truth value of (2) does not depend on whether or not our world is missing the history, h_ω ; (2) holds whether or not we count h_ω as a history. The truth value of the following, however, depends on precisely that.

At m_0 , it is inevitable that a will decay after a finite number of ticks. (3)

This may be restated in various ways.

- (i) At m_0 , it is inevitable that, sooner or later, the atom, a , decays. (ii) At m_0 , it is inevitable that the chain of no-decays terminates. (iii) It cannot be that a never decays. (iv) That a never decays is impossible. (v) At m_0 , it is false that a might never decay.

In (3) and in all of its restatements, the modal words are to be taken as historical modalities, equivalent to quantification over histories. So it matters whether the maximal chain of moments, h_ω , which is determined by the rightmost sequence of moments, counts as a history.

On our no-missing-history account, h_ω has to be a history merely in virtue of being a maximal chain in the tree: $\langle Tree, \leq \rangle$ has the histories that it has, and if $\langle Tree, \leq \rangle$ represents our world with its causal order, there is nothing more to say. Therefore the no-missing-histories account says that (3) is, in all its versions, false. Suppose, however, that h_ω does not count as a history, as is certainly allowed if a missing-histories structure, $\langle Tree, \leq, History \rangle$, can describe our world. Then all the versions of (3) are true. If you leave out h_ω as a history, then no matter what history, the sequence of no-decay terminates after a finite number of ticks.

Does this quarrel finish in a draw between the missing-histories and no-missing-histories representations of our world? We don't think so. It seems to us plain, following an analogous verdict by Thomason 1984, pp. 151–152 (but contrary to the verdict of Øhrstrøm and Hasle 1995, pp. 268–269) that anyone who asserts both (2) and (3) has contradicted himself. Surely, we say, it is a real possibility, not to be ruled out by switching “logic,” that the atom may never decay.

There is also an argument against the legitimacy of the missing-histories representation that does not depend on intuitions concerning the validity of hard-to-understand arguments. We introduced moments as representing concrete possible events, and \leq as representing the (indeterministic) causal ordering among them. This gives us $\langle Tree, \leq \rangle$ as rooted in objective reality—idealized, of course. The set of *all* histories is uniquely determined in terms of *Tree* and \leq , so that we take that set as itself objective rather than made up to suit conversational context, or language, or the like. Let us now consider a missing-histories structure, $\langle Tree, \leq, History \rangle$, where one or more histories is missing from *History*. How can we see as *objective* the separation of histories (maximal chains) into those that belong to *History* and those that do not? Using “chronicle” where we say “history,” Øhrstrøm and Hasle 1995 argue that we *must* “assume that not all linear subsets [of *Tree*] are possible chronicles.” But what objective property of our world could justify treating some maximal chains as real possibilities and others as not? These questions seem to us to have only implausible answers; consult the following observations.

Observations. (i) Probabilities don't come into it, just possibilities. One may wonder, however, if a case could be made for missing histories by forging a conceptual identification of “impossible” with “zero probability.” No, for standard reasons: It is all too likely that in our world, *every* endless branch has zero probability, but there is no sanity in letting *all* of them go missing. (ii) It is no good subtracting from a tree *all* those branches containing only points that belong to other branches. It is indeed true that there is only one of these in Figure 7.2; it is, however, all too likely that in our world, every branch contains only points that belong to other branches. (iii) Some persons might think that Omnipotent God can rig things so that both (2) and (3) are




Figure 7.3: Missing history in $T \times W$ perspective

true. God could arrange things, for instance, so that on the one hand, as long as the atom has not decayed, it is guaranteed that it has another chance not to decay, and nevertheless, God forbids that it should never decay. Or maybe lawful nature can rig things in this way. We don't understand it. How could even God, or nature, prohibit that the atom will never decay, given that each stage of nondecay can be prolonged? You can *say* it of course, but does it make sense? We doubt it.

Coming back to an earlier point, each of the bundled tree and the Ockhamist structure and the Kamp-structure representations, being equivalent to the moment-history representation, deserves the same missing-histories complaint, as does $T \times W$. In the case of $T \times W$ and Kamp structures, however, it is considerably less obvious that there is an unreasonable "omission" rather than a reasonable "resistance to addition" (A. Zanardo, correspondence). Let us spell this out. In, for example, $T \times W$, one might have the infinite collection of *disjoint* worlds $\{w_1, \dots, w_5, \dots\}$, as in Figure 7.3. In this figure we represent a single moment by means of a collection of points, one for each world, that sit on the same level. A single point then corresponds not to a moment, but to a moment-history pair. The heavy dots represent branch points, and the topmost dots, marked "*", represent the last moment of no-decay in that world. The $T \times W$ diagram of Figure 7.3, with or without w_ω , is representationally equivalent to the branching-time diagram of Figure 7.2, respectively with or without h_ω . The point is that if you stare at just the infinite collection of worlds $\{w_0, \dots\}$, without w_ω , you may well *not* feel that there is a "missing world." You may instead feel that in passing from $\{w_0, \dots\}$ to $\{w_0, \dots; w_\omega\}$, an unexpected world has been *added*.

We certainly think that such a feeling is justified by the diagram. From that perspective, it seems an open question as to whether "limit worlds" such as w_ω should or should not be added in all cases. You may share in the feeling that

when *History* is the set of *all* histories, $\langle Tree, \leq, History \rangle$ is somehow unusual or special. You may therefore prefer to abandon the language of “missing histories” versus “no missing histories” that we have employed because it suggests that having all the histories is philosophically normal. You may instead prefer to speak of “moment-history structure” versus “*complete* moment-history structure,” with its suggestion that having all the histories is special. Are we then back to seeing the quarrel as finishing in a draw? We don’t think so. We think that the $T \times W$ picture is a *mere* diagrammatic representation: There is more in the diagram than there is in our world. In particular, there seems to us no objective truth to all of those disjoint “worlds” in Figure 7.3 except to the extent that they jointly represent our one objectively real world with its concrete events in their indeterministic causal order giving rise to a system of overlapping, branching histories. We recommend not trusting diagrams like Figure 7.3 when they are not rooted in objective features of our only world. To put the matter another way, if what binds all the points on a certain level into the representation of a single concrete event is not sheer identity, then there is nothing else objective for it to be. “Matching” is a myth.

One may take the fact that $T \times W$ or Kamp diagrams can mislead as an additional complaint against them: They conceal the truth about the structure of our world by means of too many henscratches and too much loose play in their free-floating *History* parameter. On the other hand, we certainly recognize that to the extent that we are arguing from premisses, our argument is circular. Someone who believes that the $T \times W$ diagram of Figure 7.3 gives the right picture of the “facts” of the decay example of Figure 7.2 will draw quite the opposite conclusion about both.

7B Theoretical reflections on indeterminism

Before passing on to consider the postulates governing agents and their choices, we sharpen our understanding of determinism/indeterminism, and clarify the idea of branching histories by considering some objections.

7B.1 Determinism and its denial

This theory about *Tree* does not contain an explicit denial of determinism, but at least we can give a slightly freshened account of what determinism means. You will note that the account involves neither laws nor theories nor any other human creations. It is in this sense “objective” or “natural.” First we need a concept.

UNDIVIDED. (*Definition. Reference: Def. 4*) $h_1 \equiv_{m_0} h_2$ iff $m_0 \in h_1 \cap h_2$ and there is an m_1 such that $m_0 < m_1$, and $m_1 \in h_1 \cap h_2$ (unless there is no m_1 such that $m_0 < m_1$). We say that h_1 and h_2 are *undivided at* m_0 . We adapt undividedness-at to pasts as well: Two histories extending a past, p , are *undivided at* p iff they share a moment properly later than p , so that they appear as a single line as p comes to a close.

That is, not only must the histories share m or p , but they must also share some properly later moment. Now we can say with absolute clarity what it is for a concrete situation to be deterministic.

DETERMINISTIC. (*Definition. Reference: Def. 4*) *Tree* is deterministic at a moment, m [or at a past, p] iff every pair of histories through m [or p] is undivided at m [or p].

There may be many histories through a deterministic past, p , but if so, they must become “many” after p is past; there is no branching as p itself comes to a close if p is deterministic. There is another way of saying this that involves the notion of *immediate* possibility.

IMMEDIATE POSSIBILITY. (*Definition. Reference: Def. 4*) A set, H , of histories is an *immediate possibility* at a moment, m_0 iff H is a subset of $H_{(m_0)}$ that is closed under undividedness at m_0 : $(h_1 \in H \text{ and } h_1 \equiv_{m_0} h_2) \rightarrow h_2 \in H$.

So for *Tree* to be deterministic at m_0 is for there to be but a single *immediate* possibility at m_0 . Sounds right.

Obviously *Tree* may be deterministic at m but not so at either earlier or later moments. On the present account of determinism, one can coherently believe that our world used to be deterministic but is not so now, although it may become so once again. We go on to say that *Tree* itself is deterministic if it is deterministic at every past. In that case there is obviously but a single history. A *determinist* is someone who believes that the tree (world) of which our moment is a part is deterministic. It would appear that many philosophers believe that anyone who is not a determinist is softheaded and probably needs therapy. Others believe that anyhow all respectable philosophical theories, including theories of agency, should at least be consistent with determinism. *Determinism, however, is an extremely strong theory*, going far beyond determinism of the present moment. In any event, we are not determinists, even though the denial of determinism is not a postulate of this book. But more than that, on the theory here offered, if anyone could ever see to anything, then determinism is false. So even though we do not lay down indeterminism as a postulate, since we believe that sometimes people have choices, we are indeterminists. Accordingly we think that any theory (of anything) should be compatible with at least a little indeterminism. We are “compatibilists” in the best sense. We agree with Kane 1998 that in particular the question “whether a kind of freedom that requires *indeterminism* can be made intelligible” (p. 105) deserves, instead of a superficial negative, our most serious attention, and indeed we intend that this book contribute to what Kane calls “the intelligibility question.”

Note, incidentally, that some situations in our world could, for all we know, be governed by indeterminist laws that are nonprobabilistic. Such a law might describe the sorts of possible outcomes for some type of initial event, but without carrying information concerning the relative probability of those outcomes. Agency may or may not be like that.

It is abstractly interesting to observe that our world could be *strongly antiderministic* in the sense that there is splitting of histories at every moment, so that $H_{(m_1)} = H_{(m_2)}$ implies $m_1 = m_2$ for every m_1 and m_2 in *Tree*. It seems to us entirely possible that strong antiderminism is true—but maybe not. Certainly it would be mathematically attractive to be able sometimes to think of a moment as a set of histories (those that pass through it), just as we can think of a history as a set of moments (those that it passes through). With equal certainty, however, this is not a good place at which to argue from beauty to truth.

As a final word on this topic, we note that even “strong” antiderminism is not all that strong. Even “strong” antiderminism is consistent with the truth of numerous and important determinist theories about “systems” and “states” and such. Indeterminism is not disorder.

7B.2 Arguments against branching

In a passage that we also quote on p. 170, Lewis 1986 advances persuasive arguments against branching time.

The trouble with branching exactly is that it conflicts with our ordinary presupposition that we have a single future. If two futures are equally mine, one with a sea fight tomorrow and one without, it is nonsense to wonder which way it will be—it will be both ways—and yet I do wonder. The theory of branching suits those who think this wondering *is* nonsense. Or those who think the wondering makes sense only if reconstrued: You have leave to wonder about the sea fight, provided that really you wonder not about what tomorrow will bring but about what today predetermines. (pp. 207–208 of Lewis 1986; we quote the remainder of this passage later, on p. 208)

In addition to our earlier comments, we respond to this argument in two quite different ways. The first response is that it takes our ordinary ways of thinking too seriously. We draw an analogy, due to Burgess 1978 (p. 165), between (i) objections to “some of our ordinary ways of speaking” based on standard relativity theory and (ii) objections to some of those ways based on branching time. In fact we sharpen the Burgess analogy by emphasizing the similarity of the roles of “many histories” in indeterminism and of “many frames of reference” in the theory of relativity. When taken as a piece of argumentation, the analogy could go like this. If the foundation of the Lewis argument in our ordinary ways were solid, then the following would be an easy reduction to absurdity of the theory of relativity.

The trouble with the theory of relativity exactly is that it conflicts with our ordinary presupposition that we have a single present or “now.” If two presents or “nows” are equally mine, one with a sea fight at Neptune’s north pole and one without, it is nonsense to wonder which way it now is—it is now both ways—and yet we do wonder.

That is a bad argument. Certainly many of us have a tendency to wonder what is going on at Neptune's north pole (or on the far side of Sagittarius) right "now." The presupposition underlying this wondering is "ordinary" and perhaps even "natural" for clock-aware and scientifically sophisticated persons born in the last few hundred years. But nevertheless it is false, as we are taught by relativity. We urge as apt the analogy suggested by this rewording of the quoted argument against branching time: Relativity insists that our world provides us with no uniquely natural spatially extended "now," although it may permit us to consider, if we wish, the limited family of all nows ("hyperplanes") to which a given utterance-event belongs. And indeterminism insists that our world provides us with no uniquely natural spatiotemporally extended actual history, although it permits us to consider, if we wish, the limited family of all histories to which a given utterance-event belongs. To the extent that common sense asks for a unique naturally given "now" to which a given utterance-event belongs, or for a unique naturally given "actual history" to which a given utterance-event belongs, to that extent, common sense is asking for something it cannot have.

Nevertheless, the Prior-Thomason semantics, which explicitly recognizes the relativity of many statements to histories as well as moments, can give common sense a large amount of what it wants and can correct some parts of the Lewis formulation that are too hasty.

"*We have a single future.*" If this means that it is settled what will happen, for example, that either it is settled that there will be a sea battle or settled that there will not, it is false. If it means that it is settled that incompatible events will never happen, it is true. If it means that there is a single future history following upon this utterance, it is false. If it means there is a single future of possibilities, it is true.

Branching time indeed claims that it can happen that "*two futures are equally mine, one with a sea fight tomorrow and one without.*" But what does it mean to say that a future is "mine"? Branching time says that it means that it is among the futures now possible, where the "now" is indexically mine. To avoid making branching time look silly in a way that it surely isn't silly, the quoted description should be amended by insertion as follows: "*two possible futures are equally mine, one with a sea fight tomorrow and one without.*" Lewis misdescribes the theory of branching time in saying of such a situation that "*it will be both ways.*" Branching time is entirely clear that "Tomorrow there will be a sea fight and tomorrow there will not be a sea fight" is a contradiction. What is true and not surprising is that "It is possible that tomorrow there will be a sea fight and it is possible that tomorrow there will not be a sea fight" is eminently consistent.

Nor is this merely a matter of formal tense logic. It seems to us deeply realistic to take it that if the captain is faced with two possibilities, sea battle tomorrow or no sea battle tomorrow, then those are possibilities for him, on that occasion. They are *equally* his, not one more than the other, exactly in accord with Lewis's account of (not his own theory but) branching time. Suppose the sea battle comes to pass. Then (after the sea battle) the two possibilities *were* his, and *were equally* his. In particular, branching time rejects the Lewis

shadow-theory according to which the captain himself is to be found in the world of the sea battle, whereas merely one of his "counterparts" can be located in worlds without sea battles. The point goes back at least to Burgess 1978, who reflects as follows on his earlier decision to go to the office instead of the seashore: "Tense logic insists, *pace* Lewis, that I am the very same person who could have gone to the shore; it's not just someone like me who could have gone" (p. 173).

Suppose we "step outside of branching time." To do this is to confine ourselves to language that has no trace of indexicality, a perfectly proper thing to do (see note 14 on p. 162). Histories and moments and persons are then linguistically accessible only via (rigid) naming or quantification. Branching time then says that for suitable moment m_0 and histories h_1 and h_2 , the captain lives through a sea battle the day after m_0 on h_1 , and lives through no-sea-battle the day after m_0 on h_2 . Here seems a premiss for a *reductio* of branching time, for branching time then seems to say that the captain has it both ways, both living through a sea battle and living through no-sea-battle.

The *reductio* is, however, an illusion. Omitting the relativization to histories is intolerable. What branching time says is that the captain "has it both ways" in the entirely innocuous sense that he lives through a sea battle on history h_1 and lives through no-sea-battle on history h_2 . That just says that there are at m_0 two possibilities for him, a fact about our world that we must keep. It does not say that the two possibilities will each be realized, an absurdity that branching time denies. It does not say that these possibilities remain possibilities at moments after the sea battle has commenced. It only says that in the past of such moments the two possibilities were available. Current possibilities drop off (McCall 1994) with passage into the future, but not the fact that they once were possibilities. Once was-possible, always was-possible.

What about "*wondering*" whether or not there will be a sea battle? Evidently our wondering is history-independent: The fact that we wonder is dependent on the moment but independent of the history parameter. So what sense can we make of wondering about a history-dependent complement such as "there will be a sea battle tomorrow"? Lewis points out one alternative, which he rightly presents as not very ordinary, namely, that the complement of the wondering is the history-independent question, whether or not it is now *settled* that there will be a sea battle tomorrow. We proposed in §6E a two-point understanding of wondering that lets its complement remain open in the history parameter. First, it seems natural to construe "*wondering* q ," where q is an indirect question (e.g., "whether there will be a sea battle tomorrow"), as "wanting to know" (or perhaps "wanting to have") a true answer to the question of q . Second, it seems obvious that our wants are not normally satisfied immediately; we must in general *wait*. Just so, in order to obtain satisfaction of the want expressed in our wondering, we shall need to wait until tomorrow; for only tomorrow is it possible for us to come to know whether or not a sea battle comes to pass.

We next comment on some ideas in the remainder of the passage the beginning of which we quoted once on p. 170 and again on p. 205. Lewis continues in the following way.

But a modal realist who thinks in the ordinary way that it makes sense to wonder what *the* future will bring, and who distinguishes this from wondering what is already predetermined, will reject branching in favour of divergence. In divergence also there are many futures; that is, there are many later segments of worlds that begin by duplicating initial segments of our world. But in divergence, only one of these futures is truly ours. The rest belong not to us but to our other-worldly counterparts. Our future is the one that is part of the same world as ourselves. It alone is connected to us by the relations—the (strictly or analogically) spatiotemporal relations, or perhaps natural external relations generally—that unify a world. It alone is influenced causally by what we do and how we are in the present. We wonder which one is the future that has a special relation to ourselves. We care about it in a way that we do not care about all the other-worldly futures. Branching, and the limited overlap it requires, are to be rejected as making nonsense of the way we take ourselves to be related to our futures; and divergence without overlap is to be preferred. (p. 208)

The first sentence refers to “*the* future,” the one that “is truly ours.” Branching time says that only the future of possibilities is uniquely determined by the moment of utterance, so that “the future” either refers to this, or else is not history-independent (is open in the history parameter). Branching time says that if indeterminism be true, then there is no more sense to “the actual future” than there is to “the actual distant instantaneous present” or to “the odd prime number.”

But what about the future that “*is part of the same world as ourselves*”? Assuming indeterminism, there is the following dilemma.

- If we read “world” as “history,” then it makes no sense to speak of “*the* world of which we are part.” There are many such possible histories to which this utterance-event equally belongs. All of them are “connected to us by the ... spatiotemporal relations ... that unify a world,” for there is, in our opinion, no more fundamental “natural external relation” than the causal ordering itself. It is to be borne in mind that even wholly incompatible moments are mediately connected by \leq ; that is exactly the import of historical connection. It is why (or how) *Tree* constitutes a single world, our world.
- If we read “world” as *Tree* in its entirety, then although it would make sense to speak of “*the* world of which we are part,” it would not make sense to speak of “*the* future” history that is part of that world. On the other hand, invoking a familiar contrast, it would indeed make sense to ponder “*the* future of possibilities.”

Lewis then gives what are in effect three arguments that, contrary to branching time, “this very event” picks out a unique future history.

- The first argument is that we can define “the future history” as the one that “alone is influenced causally by what we do and how we are in the present.” This sounds all right, but it is not. It is not merely that we cannot evaluate this proposal in the absence of an objective theory of “causal influence” and of “what we do.” Nor is it just a matter that the argument won’t tell against a theory such as ours that *explicitly* holds that if there is only one possible future history issuing from this present event, then, the future being settled, there is no “influencing” the future by what we do in the present. The difficulty is above all that the proposed definition of “the future history” is akin to defining “the odd prime number” as the one that “alone is not divisible by two.”
- The second argument seems to be that *the* future is the one we wonder about. As we indicated, however, at the very end of chapter 6, and again a couple of pages ago, to the extent that wondering is wanting to know, wondering is similar to other wants. Whether one wants to know what will happen as the open future becomes determinate, or one wants to have a dapple gray pony, one must bide one’s time.
- The third argument is about *caring*. Certainly branching time shares with this passage the premiss that we do not care about other-worldly futures. Indeed the passage suggests a certain mild *ad hominem*: What we care about is what alternatives there are *for us* to choose among. We do not (much) care what alternatives there are for other-worldly counterparts, should they exist as Lewis’s theory requires. We don’t see how counterpart theory can either make sense of or reject the demand that future (incompatible) possibilities be *for us* rather than for our counterparts.

Suppose for example that we are choosing whether or not to start a sea battle. Surely we care about what will happen if there is a sea battle, and we also care about what will happen if there is not a sea battle, since these are possibilities for us. And given that we really do have a choice, and know that we do, such caring about incompatible alternatives makes sense. If these histories are (right now) all really possible, then we do now rightly care what is true on each. We really do care about what happens on more than one history—as long as the histories are ours.

7B.3 Cusp of causality

There is more than one theory of causality in branching time; see von Kutschera 1993 and Xu 1997. Here we ask: Suppose that *A*, which in general is dependent on the history of evaluation, is nevertheless settled true at *m*. *A* has become a settled “fact.” *When* did this become so? If we can answer this, we shall have found a causal locus for the “effect” that *A* is settled true at *m*. If *A* is the sort of sentence that changes its truth value over time, then the question may be difficult (Belnap 1996a gives an unsatisfying answer). If, however, the truth of *A* is moment-independent, then the question becomes manageable. In this case,

A is variable only in the history parameter. "Dated" sentences such as "At 4:00 P.M., A ," or renditions of "The coin will land heads sometime tomorrow" have this feature: Their truth depends on the history of evaluation (and in the second example also on the moment of use), but not on the moment of evaluation.

So suppose A is moment-independent. And suppose that A is in fact settled true at m . To look for a causal locus, we rely on the following.

7-1 FACT. (*Settledness of moment-independent sentences*) If the truth of A is independent of the moment parameter, then settledness propagates forward and lack of settledness (openness) propagates backward: If $m_1 < m_2$, then if A is settled true [false] at m_1 , it is also settled true [false] at m_2 ; and if A is not settled true [false] at m_2 , it is also not settled true [false] at m_1 .

Consider now the improper past of m , call it p . It could be that A has been settled true throughout p , in which case, since settledness propagates forward, A is "universally true," that is, A is settled true throughout all of branching time. Suppose, then, that m has not been settled true from all eternity. Then because of Fact 7-1, we may make a Dedekind cut of p into two nonempty chains: At every moment in the lower portion, call it c_{open} , the truth of A will depend on the history parameter; while at every moment in the upper portion, call it c_{sett} , A will be settled true. Thus, where c_{open} draws to a close is the very "point" at which the status of A changes from "not yet settled to be true" to "now settled to be true." We put "point" in shudder-quotes because c_{open} may not end in a moment, so for definiteness we call all of c_{open} *the cusp of causality*.

The cusp of causality is where to look for the causal locus of the "effect" that A is settled true at m . It is well to keep in mind that m , where we have supposed A to be settled true, may itself minimally upper bound the cusp of causality. For example, if the A in question is that we are at the restaurant at 6:00 P.M., Murphy's Law suggests that it may well take right smack up to 6:00 P.M. in order to settle that fact in our favor; in which case the cusp of causality is the entire set of proper predecessors of the moment of our being at the restaurant at 6:00 P.M.. Some people think this is a defect in the theory; we think it is a defect in the world.

Though certainly imperfect in that it is nonrelativistic, the concept of the cusp of causality is nevertheless an empirical, objective causal concept. The idea is that nothing can be an effect unless it changes from not-being-settled-to-be-true to being-settled-to-be-true, and that the locus of causality must be at the "point" at which this change takes place: The cusp of causality is just where the effect comes to be settled true.

7C Theory of agents and choices

We have finished discussing the postulates governing the $BT + I$ part ($\langle Tree, \leq, Instant \rangle$) of $BT + I + AC$ structures ($\langle Tree, \leq, Instant, Agent, Choice, Domain \rangle$), that is, the part of those structures that represent indeterminism and add the

possibility of making cross-history time comparisons. In this section we go on to consider the remaining parameters *Agent* and *Choice*, with the aim of characterizing how agents and their choices fit into the causal structure of our world as represented by $\langle Tree, \leq, Instant \rangle$, and in the next section we add a coda about *Domain*.

7C.1 Agents

Nowhere in this book do we offer thoughts that help much when considering the mental equipment or other aspects of the “real internal constitution” of agents. The topic is important, and important to us; it is just that our attempt at progress in this book takes us in a different direction. Alan Ross Anderson used to say that all progress in philosophy comes by simply assuming certain problems have been solved even though they haven’t, and getting on with the investigation. We pursue this policy by entering a significantly uninformative postulate concerning agents.

AGENTS. (*BT + I + AC* postulate. Reference: **Post. 6**) *Agent* is a nonempty set.

We call the members of *Agent* *agents*, and we let lowercase Greek letters α and β range over agents. We intend that the concept of *Agent* is absolute in the sense of Bressan 1972 (or a substance sort in the sense of Gupta 1980), which means that we may “identify” agents across times and histories. In particular, there is no fission and no fusion of agents as we move from moment to moment. In this book, however, we do not happen to discuss questions such as *de re* versus *de dicto* ascriptions of agency, nor do we worry about when agents come to be or pass away. That is why we can get by with the simpler set-theoretical representation of *Agent* given by **Post. 6**. There is a brief discussion of *Agent* as an absolute concept in §10C.1, where we are worrying about joint agency, and a little more in §12F, where we are thinking about generalizing on the agent position of stit statements. None of this, however, is offered as a serious contribution to solving the problem of personal identity. On the other hand, we do think that consideration of agents in branching time may deepen that problem. *BT + AC* theory, for example when considering strategies, characterizes the *same* agent as making some (possible) choices sequentially, and also some (possible) choices under incompatible circumstances. That the idea of “same agent” seems essential to the idea of, for example, a strategy is, we think, a good reason to suppose that a concept of personal identity that does not *essentially* involve agency is much too partial. A concept of personal identity that depends on exclusively “passive” notions such as experiential content, or on only backward-looking ideas such as memory, is by so much inadequate.

7C.2 Choices

The penultimate parameter of a *BT + I + AC* structure is *Choice*, which tells what choices are open for each agent at each moment in *Tree*.

CHOICE PARTITION. (*BT + I + AC postulate. Reference: Post. 7*) *Choice* is a function defined on agents and moments. Its value for agent α and moment m is written as $Choice_m^\alpha$. For each agent α and moment m , $Choice_m^\alpha$ is a partition into equivalence classes of the set $H_{(m)}$ of all histories to which m belongs.

The postulate may strike one as unhappy to the extent that it seems to insist that choosing is localized in a moment, and is therefore instantaneous. The opposite seems so natural: Choosing takes time. We therefore feel obligated to explain, in part, our reasons for giving a key role in *BT + I + AC* theory to such a *prima facie* counter-intuitive postulate.

Jack has made a deliberate choice to go to the beach. There is, we think, no harm in postulating that there is a momentary event that entirely precedes his beginning to deliberate; let it be m_1 . There is also no harm in postulating a momentary event that lies thoroughly after his arrival at the beach; let it be m_2 . To find a plausible m_1 and m_2 we don't need to know *exactly* how Jack's deliberation relates to its outcome; we only need to know that his deliberation has a beginning that precedes his arrival at the beach. Having fixed m_1 and m_2 in this way, if deliberation is to have point, it must not be decided at m_1 , before the deliberation begins, whether or not Jack later arrives at the beach (Aristotle). And at m_2 , after Jack arrives, it must certainly be decided that he arrives (how else?). So now draw a chain of moments from m_1 to m_2 . Since that chain has members, which we call moments, then by simple Dedekind analysis, there must either be a last moment of undecidedness or a first moment of decidedness, or both (a Dedekind jump) or neither (a Dedekind gap). Our theory of branching time does not say which; it only says that there is a *transition* from undecidedness to decidedness. Nevertheless, our thoughts go—we think harmlessly—with those who think of the flow of our world as continuous, and hence without jumps or gaps. Furthermore, it is certain that if our world is continuous, and can be accurately represented by a continuum of instantaneous moments, then there must be either a last instantaneous moment of undecidedness or a first such of decidedness. Since epistemology is not likely to help us choose between these two, we see no objection to our always thinking of a last instantaneous moment of undecidedness, a last moment at which it is still not decided whether Jack will arrive at the beach. (Evidently there will also be such a last moment of undecidedness if our world proceeds by discrete jumps.) Talk of “deliberation flowing into action” is all right unless it blinds us to these observations, and especially to the primacy of an objectively more-or-less localizable *transition* from undecidedness into decidedness. It is this idea that is idealized in **Post. 7**. We theoretically identify “the moment of choice” as the last moment before the matter is decided, while still thinking of choice itself as fundamentally a *transition* from undecidedness to decidedness. (See §2A.2 for discussion.)

Observe that nothing in the postulate denies that deliberating takes time. If deliberation is mere wheel-spinning, we say nothing about it. If, however, it involves, for example, a continuous (or discrete) ruling-out of alternatives, we should then represent it theoretically as a continuum (or succession) of choices.

Finally, we observe that the postulate says nothing about whether choice is localized in some homunculus in the brain, or whether it is always made by a neuronal "community," or indeed whether choice is in any way localized. Nor does any part of our theory say anything about what choice feels like, or whether every choice is a conscious choice. These are difficult and important questions. We remind the reader here, however, as we do elsewhere, that our explicitly and advisedly limited theory concerns only the causal structure of choice, to the exclusion of its "content." Nor will it hurt to insert our view that those who deal with the "content" of choice to the exclusion of its (indeterminist) causal structure can easily be led astray.

The fundamental postulate **Post. 7** on choice says that histories are divided evenly into equivalence classes; we discuss this aspect after we introduce the equivalence-relation notation for choices, together with some closely related notation that we use with considerable frequency.

CHOICE NOTATION. (*Definition. Reference: Def. 11*)

- i. *Choice* represents all the choice information for the entire *Tree*.
- ii. $Choice_m^\alpha$ gives all the choice information for the agent α and the moment m ; $Choice_m^\alpha$ should be thought of as a set of possible choices, and we call it "the set of choices possible for α at m ."

By **Post. 7** we know that $Choice_m^\alpha$ is a partition of $H_{(m)}$: Each history on which m lies belongs to exactly one member of $Choice_m^\alpha$. It is essential to keep in mind that a possible choice is a *set* of histories, not a single history. No one—of course—can choose their future down to the smallest detail.
- iii. $Choice_m^\alpha(h)$ is defined only when h passes through m , and is then the unique possible choice (a set of histories) for α at m to which h belongs. The notation is justified by the fact that according to **Post. 7**, each member of $H_{(m)}$ picks out a unique member of $Choice_m^\alpha$ to which it belongs.
- iv. $Choice_m^\alpha(m_1)$ is defined only when m_1 is in the proper future of the moment m of choice. Pick any history, h , containing m_1 , a history that will *a fortiori* contain m . Then $Choice_m^\alpha(m_1)$ is defined as $Choice_m^\alpha(h)$. This definition is justified by no choice between undivided histories, **Post. 8**, which we have not yet discussed in this chapter. We include the notation in this list for reference. Its theoretical import is this: If we are at a moment m_1 , such that a moment of choice, m , lies in its past, the "the choice that α made at m " is uniquely determined; it is a settled fact.
- v. $\underline{Choice}_m^\alpha(m_1)$ is defined only when instants are present, and when m_1 is properly future to m . Recall that $i_{(m_1)}$ is the instant on which m_1 lies. Then $\underline{Choice}_m^\alpha(m_1)$ is defined as the set of all moments on the instant $i_{(m_1)}$ that also lie on some history in $Choice_m^\alpha(m_1)$. In symbols, $\underline{Choice}_m^\alpha(m_1) = i_{(m_1)} \cap \bigcup Choice_m^\alpha(m_1)$. If we are at a moment m_1 , such that a moment

of choice, m , lies in the past, then $\underline{Choice}_m^\alpha(m_1)$ picks out the “projection” of the uniquely determined choice made by α at m onto the present instant or time.⁶

It is the intent of the postulate **Post. 7** on *Choice* that when agent and moment are fixed, choice equivalence is an equivalence relation on the histories passing through the moment: reflexive, symmetric, and transitive. For example, to say that $Choice_m^\alpha(h_1) = Choice_m^\alpha(h_2)$ is another way of saying that h_1 belongs to $Choice_m^\alpha(h_2)$. The next definition describes the equivalence-relation notation for choices.

CHOICE EQUIVALENCE. (*Definition. Reference: Def. 12*)

- $h_1 \equiv_m^\alpha h_2$ iff $Choice_m^\alpha(h_1) = Choice_m^\alpha(h_2)$, and we say that h_1 and h_2 are *choice equivalent* for α at m .
- $h_1 \perp_m^\alpha h_2$ iff $m \in h_1 \cap h_2$ and $Choice_m^\alpha(h_1) \neq Choice_m^\alpha(h_2)$. We say that h_1 and h_2 are *choice separated* for α at m .
- $m_1 \equiv_m^\alpha m_2$ is defined only when instants are present, and when $m < m_1, m_2$. Then $m_1 \equiv_m^\alpha m_2$ iff $\underline{Choice}_m^\alpha(m_1) = \underline{Choice}_m^\alpha(m_2)$. We say that m_1 is *choice equivalent* to m_2 for α at m .

$h_1 \perp_m^\alpha h_2$ is contrary to $h_1 \equiv_m^\alpha h_2$. It makes sense to think of a given possible choice as “separating” each history in the choice from each history not in the choice. So if we wish to describe separation in a sentence that makes “ α ” the grammatical subject, it is all right to say something like “ α has a possible choice at m_0 that separates h_1 and h_2 .” We should not give in to the temptation to say that “ α can choose between h_1 and h_2 .” To put the matter as clearly and therefore trivially as possible, α can never choose between histories, but can only choose between choices.

We now consider whether *Choice* is correctly described as a partition, or, what is technically equivalent, whether choice equivalence for α at m is a reflexive, symmetric, and transitive relation on $H_{(m)}$.

Reflexivity is perhaps a “throwaway” postulate, except to the extent that it implies that *every* history belongs to some possible choice. Perhaps it makes sense to say that at some moment, if certain things happen α has chosen, but if other things happen, α has not chosen; J. MacFarlane has suggested some potential cases, and has urged that we keep our minds open to this possibility. Nevertheless, it seems to us that speaking causally, and regardless of how people evaluate a given agent or situation, we tend to side with those who find conceptual difficulty in an alleged situation involving the simultaneous possibility of choice and no-choice. We do, however, agree with MacFarlane that we know much too little to be warranted in discouraging the development of alternative theories along these lines.

⁶ $Choice_m^\alpha(m_1)$ gives a set of *histories* and $\underline{Choice}_m^\alpha(m_1)$ gives a set of co-instantial *moments*, which is confusing. The underlining on \underline{Choice} is a mnemonic intended to help keep the two concepts apart by calling to mind the horizontal picture of an instant.

Symmetry seems to us inescapable when it comes to choices.

Transitivity, as usual, is the easiest feature of a partition-postulate to doubt. Given that we have no choices that distinguish h_1 from h_2 , and no choices that distinguish h_2 from h_3 , are we forced to conclude that there are no choices for the agent that distinguish h_1 from h_3 ? This seems to us at least a quasi-empirical question, not decidable by easy philosophical speculation. Is the nontransitivity of “not noticeably different” (e.g., for colors) a reason for doubting the transitivity of choice equivalence? Whether yes or no, it would be good to have a theory that did not presuppose transitivity of choice equivalence. We have none to suggest, and so are content to “admit” that in the end our ideas, which make heavy use of transitivity, only apply when transitivity either seems plausible, or at least seems to be an idealization that does not interfere with our modest progress in understanding agency.

Especially for the theory of strategies, it is good to generalize choice equivalence and choice separability from concepts based on a single moment to concepts based on a set of moments. Think of M in these definitions as a collection of choice points lying in your future. Which histories can you tell apart by choices at these points? (See **Def. 12** for the choice-equivalence notation used in defining these ideas.)

INSEPARABILITY/SEPARABILITY. (*Definition. Reference: Def. 13*)

- $h_1 \equiv_M^\alpha h_2$, read “ h_1 is inseparable from h_2 for α in M ,” iff $\forall m[m \in M \cap h_1 \cap h_2 \rightarrow h_1 \equiv_m^\alpha h_2]$.
- $m_1 \equiv_M^\alpha m_2$, read “ m_1 is inseparable from m_2 for α in M ,” iff $\forall m[(m \in M \ \& \ m < m_1 \ \& \ m < m_2) \rightarrow m_1 \equiv_m^\alpha m_2]$.
- $h_1 \perp_M^\alpha h_2$, read “ h_1 is separable from h_2 for α in M ,” iff $\exists m[m \in M \cap h_1 \cap h_2 \ \& \ h_1 \perp_m^\alpha h_2]$.
- $m_1 \perp_M^\alpha m_2$, read “ m_1 is separable from m_2 for α in M ,” iff $\exists m[m \in M \ \& \ m < m_1 \ \& \ m < m_2 \ \& \ m_1 \perp_m^\alpha m_2]$.

The fundamental idea is carried by $h_1 \equiv_M^\alpha h_2$, which says that no choice for α in M separates h_1 and h_2 . “Inseparability” in this context is a geometrically suggestive rhetorical variation on “choice equivalence” that correctly intimates the nontransitivity of the relation. Its contrary is also conceptually important; two histories are “separable for α in M ,” for instance, if somewhere in M there is a choice for α that keeps one of the histories as a possibility while ruling out the other as thereafter impossible.

Figure 7.4 provides an example of nontransitivity of inseparability, using a simple case of a chain, c , consisting of the two moments $\{w_0, w_1\}$ represented by divided rectangles. The division of each rectangle represents that at each moment there are two possible choices for α . The history, h , that splits from the chain at the lower moment is choice inseparable for α by c from each of the histories h_1 and h_2 that split at the top moment; but obviously h_1 and h_2 are not choice inseparable for α by c from each other. That is, no choice that




Figure 7.4: Nontransitivity of choice inseparability at a chain

α makes in the course of c separates h from h_1 , or h from h_2 , but there is evidently a choice (at the top moment) that separates h_1 from h_2 . So choice inseparability for α at a chain is not and should not be transitive.

There are two further postulates relating to choices by agents. The first amounts to a “new” principle relating agency to the causal order, while the second concerns multiple or joint agency.

7C.3 No choice between undivided histories

This postulate requires the previously given temporal-modal definition of “undivided,” **Def. 4(iv)**.

NO CHOICE BETWEEN UNDIVIDED HISTORIES. (*BT+I+AC postulate. Reference: Post. 8*) If two histories are undivided at m , then no possible choice for any agent at m distinguishes between the two histories. That is, one of two histories undivided at m belongs to a certain choice possible for α at m if and only if the other belongs to exactly the same possible choice. In symbols from **Def. 4** and **Def. 12**: $h_1 \equiv_{m_0} h_2 \rightarrow h_1 \equiv_{m_0}^{\alpha} h_2$.

As reported in, for example, chapter 1, we learned the no choice between undivided histories condition from P. Kremer in 1987. All the postulates having to do with choices are to be found in one form or another even earlier in von Kutschera 1986.

This postulate is perhaps the most “interesting” of the *BT+I+AC* postulates. As far as we know, the idea of relating choice to brute causal undividedness has no earlier rigorous expression, even though the relation is and must be of importance to the theory of action and to moral theory.

An easy consequence is that from the point of view of a properly later moment m_1 , what choice an agent made at each properly earlier moment m , is uniquely determined. Or to say the same thing from the point of view of m , for each moment m_1 , properly later than m , there is a unique possible choice for α at m that contains all histories passing through m_1 . For argument, assume m_1 is properly later than m . Then the two moments constitute a (two-member) chain that by Zorn’s lemma can be extended to a maximal chain, that is, a

history, h , that contains them both. Existence is partly given by $Choice_m^\alpha(h)$, which evidently contains a history, namely, h itself, that passes through both m and m_1 . To continue with existence, let h' be any history through m_1 , and hence, by no backward branching, through m . By no choice between undivided histories, h' must belong to $Choice_m^\alpha(h)$, which then contains every history through m_1 . As for uniqueness, if any possible choice for α at m contained all histories through m_1 , it would also contain h , which would imply its identity with $Choice_m^\alpha(h)$, as required.

Provided m_1 is properly later than m , this fact justifies introducing the concept, "the choice that was made by α at m ," where the past tense is issued from the point of view of m_1 . In Def. 11(iv) we introduced the notation $Choice_m^\alpha(m_1)$ for this concept.

It follows from no choice between undivided histories that, to the extent a choice objectively limits the future, you cannot today "choose to choose" to have potatoes tomorrow. You can say "Tomorrow I will choose to have potatoes" and you can intend today that tomorrow you choose to have potatoes, and you can stay well within the bounds of commonsense language by saying "I hereby choose to have potatoes tomorrow," even when you know full well that today's choice does not determine the matter, since tomorrow you can always change your mind. You may even be in a position to choose today that you, without further choice, have potatoes tomorrow. (Perhaps you put yourself on automatic pilot.) What you cannot do is both choose today and also choose tomorrow. You can have it both ways in conversation or in your mind or in your philosophical gloss, but you cannot have it both ways in reality. One way of putting the matter is this: Adopting a plan for the future is not executing that plan; and this fact derives from the fundamental principle that there can be no choice between undivided histories. A final thought: This fundamental principle surely has consequences for the concept of "weakness of the will." It does not speak to the mental side of that concept, but it does help to place weakness of the will in the proper causal context. There would be no weakness of the will if one could choose between undivided histories.

7C.4 Independence of agents

INDEPENDENCE OF AGENTS. (*BT+I+AC postulate. Reference: Post. 9*) If there are multiple agents: For each moment and for each way of selecting one possible choice for each agent, α , from among α 's set of choices at that moment, the intersection of all the possible choices selected must contain at least one history. In symbols: for each $m \in Tree$, and for each function f_m on Agent such that $f_m(\alpha) \in Choice_m^\alpha$ for all $\alpha \in Agent$, $\bigcap \{f_m(\alpha): \alpha \in Agent\} \neq \emptyset$.

Sometimes this is thought of as "independence of choices," which is a good thought. At any one moment m , the choices possible to each agent are indeed independent in the sense, for example, that any possible choice, $Choice_m^{\alpha_1}(h_1)$, that α_1 makes is consistent with any possible choice, $Choice_m^{\alpha_2}(h_2)$, that α_2 makes: The intersection of these two sets must be nonempty. Since the choices

are simultaneous, it is certainly reasonable to think of them as independent. What might be confusing about the language of independence is this: The entire set of choices, $Choice_m^{\alpha_1}$, open to α_1 at m is by no means “independent” of the entire set of choices, $Choice_m^{\alpha_2}$, open to α_2 at m . The very fact that each member of the first set must have nonempty overlap with each member of the second is itself a fierce constraint. It is easy to see, for instance, that no two agents can possibly have exactly the same possible choices at exactly the same moment (vacuous choice aside).

Fierce or not, however, we think that this postulate is banal. If there are agents whose simultaneous choices are not independent, so that the choice of one can “influence” what it is possible for the other to choose even without priority in the causal order, then we shall need to treat in the theory of agency a phenomenon just as exotic as those discovered in the land of quantum mechanics by Einstein, Podolsky, and Rosen. (This point of view can be found in the consideration of the Prisoners’ Dilemma by Green and Bicchieri 1997, where it plays an important role.) We are in effect postulating that the only way that the choices open to one agent can depend on the choices open to another agent is if the one agent’s choices lie in the causal past of those of another agent.

That almost concludes our one-by-one discussion of the $BT + I + AC$ postulates. Before briefly considering one more postulate, we describe several concepts that prove useful in various contexts.

7C.5 Vacuous choices and busy choosers

Agents choose at “choice points,” but happily not all moments are such. Although we propose no theory of the coming to be and passing away of agents, at least we can use the idea of “vacuous choice” to avoid suggesting that each agent is constantly and forever choosing: A vacuous choice (which is, incidentally, the only sort offered under determinism) is not a choice.

VACUOUS CHOICES. (*Definition. Reference: Def. 14*)

- A moment m is a *choice point* for α iff there is more than one possible choice for α at m .
- A possible choice for α at m is *vacuous* or *trivial* iff it is the only possible choice for α at m ; and is otherwise *nonvacuous*. There can only be a vacuous choice for α at m when $Choice_m^\alpha = \{H(m)\}$, in which case m itself is said to be a *trivial choice point* for α .

The idea of a busy chooser is that of an agent that makes infinitely many choices in a finite period. It is odd that although the idea is altogether peripheral to our central ideas, in technical discussions it crops up with unexpected frequency.

BUSY CHOOSERS. (*Definition. Reference: Def. 14*)

- A chain c is a *busy choice sequence* for α iff (i) c is both lower and upper bounded in *Tree*, and (ii) c is an infinite chain of (nontrivial) choice points for α .

- α is a *busy chooser* iff some c in *Tree* is a busy choice sequence for α .

We don't pretend to know if there are any busy choosers, or even if the idea makes sense. (We don't think anyone else knows either.) All we do is keep track, as well as we can, of places in the theory of agents and choices in branching time at which it makes a difference. In this regard, the following is worth noting: Sometimes what makes a difference is the existence of an unending *forward* series of nonvacuous choices, while sometimes it is the existence of an unending *backward* series.

7D Domain

The last entry in the structure $\langle \textit{Tree}, \leq, \textit{Instant}, \textit{Agent}, \textit{Choice}, \textit{Domain} \rangle$ is *Domain*, which we employ as a field of possible denotata, and as the range of individual variables.

RICHNESS OF *Domain*. (*BT+I+AC* postulate. Reference: Post. 10) The domain of quantification, *Domain*, must include *Tree*, *History*, *Instant*, and *Agent* as subsets.

That is, we postulate that moments, histories, instants, and agents shall all be among "what there is." We enter this postulate partly for some technical reasons that emerge in chapter 8, and partly to emphasize how harmless and nonparadoxical it is to include these entities in the domain of quantification. These entities are not in the least "meta-linguistic," nor does their addition make the domain "too big" to be a set. Of course *Domain* must be subject to some artificial limitation or other in order to avoid paradox, but we think that this set-theoretical (or even ontological) problem has nothing whatever to do with agency or the causal structure of our world.

8

Indexical semantics under indeterminism

Prior chapters of this book have offered bits and pieces of a description of agency in our indeterministic world. This chapter brings us to the topic of the semantics of a language taken to be used by those (ourselves) who live in such a world. Just to have a label, we sometimes call this topic either “the semantics of indeterminism” or “indeterministic semantics.”

Agency constructions we treat in §8G; until then we emphasize branching time itself. We include (i) quantificational devices, (ii) temporal constructions, (iii) historical modalities, (iv) some mixed modalities, and (v) indexicals tied to the context of use. We include all of these items because they crop up in discussions of indeterminism, sometimes in a confusing way. We say what we have to say in the idiom of formal semantics. This amounts to an organized account of the semantics needed for a language spoken in an indeterminist world. First we go over and extend some of the foundational and generic semantic ideas broached in chapter 6. Then we go one by one through a large number of constructions useful for understanding indeterminism, in each case giving an exact semantic account. We emphasize points that we take to be important, and we draw out the analogies among and differences between the semantics of quantifiers, historical-modal and tense connectives, and indexical connectives, and indicate how they are to be combined. We briefly review the already-presented semantics for the achievement stit and the deliberative stit, and we explain how stits might be witnessed by a chain of choices (instead of a single choice). The chapter ends with a mention of an alternative agency construction, the transition stit. Let us note that in compensation for the inevitable tedium of processing henscratches, the early sections expand on the fundamental ideas that go into the formal semantics.

8A Sources

As we briefly noted at the beginning of §6B, we draw on four sources for the key ideas needed for the semantics of indeterminism.

- *Quantification.* Tarski's studies in the 1930s provide the foundation for all compositional semantics, and in particular are the source for our compositional understanding of quantifiers.
- *Tenses.* Prior 1957 initiated the compositional understanding of linear tenses.
- *Historical modalities.* Prior 1967 as made rigorous by Thomason 1970 (see also Thomason 1984) is the source of our understanding of the historical modalities, and also of how tenses should be understood given a representation of indeterminism by means of branching time.
- *Context of use.* Kaplan 1989, much of which had circulated in typescript since Kaplan's 1971 lectures, provided the full-scale development of "indexical semantics" based on the idea that compositional semantics must pay delicate attention to the context of the use of an expression.

The semantics of indeterminism relies heavily on all four of these sources, which we discuss in turn as ways of following out the fundamental Tarski idea of relativizing truth to parameters.

8A.1 Tarski's quantifier semantics

Tarski saw (among much else) that, if the account of truth for quantificational sentences is to proceed in a straightforward compositional fashion, one must *relativize truth to new parameters* that go beyond an explanation or translation of the various constant features of the language (including a specification of domain). Those new parameters are *assignments of values to the variables*. Why did it take someone as exceptional as Tarski to attain this insight? Perhaps because most of us think primarily in terms of *stand-alone* sentences (vehicles of assertions), which should contain no free variables. In this case one does not need the assignment parameters. One needs them only when thinking of a sentence as arbitrarily *embeddable*, specifically including the possibility of embedding within the scope of a quantifier. And one must think of sentences in this way if one is to pass recursively from a semantic account of Fx_1x_2 to a semantic account of $\exists x_2Fx_1x_2$, and thence to an account of a potential stand-alone sentence, $\forall x_1\exists x_2Fx_1x_2$. It does not suffice to have an *absolute* "truth value" for the embedded open sentences Fx_1x_2 and $\exists x_2Fx_1x_2$. One needs to know the truth value of these embedded sentences *relative to* each appropriate family of assignments of values to the variables. In the deepest sense, a quantifier such as $\forall x_1$ is not "truth functional."

Tarski himself perhaps makes this a little difficult to see by his linguistic detour through his "satisfaction" relation, but the underlying point shines through:

Open sentences such as Fx_1x_2 do not have an absolute truth value. They rather have a truth value only *relative to* assignments of values to the variables, so that such assignments are parameters of truth.

Let us be explicit about how we extract the assignment parameters from Tarski's account so that those familiar with Tarski may understand our stylistic departure. His fundamental locution is illustrated by

$$Fx_2x_1 \text{ is satisfied by } \langle a, b, \dots \rangle, \quad (1)$$

where the members of the sequence such as a and b are individuals drawn from the domain. To make this work, one must know which member of the sequence, $\langle a, b, \dots \rangle$, goes with which variable. In this example, we have used subscripts whose sole purpose is implicitly to supply an *ordering* of the variables. This order is to be used in conjunction with the order of the sequence, so that we know that a goes with x_1 and b with x_2 . The ordering of the variables is essential to the meaningfulness of the Tarski relation of satisfaction by sequence: One must understand which member of the sequence goes with which variable. As Tarski observes, however, one does not need for this purpose to suppose a primordial ordering of the variables such as supplied by subscripts. A more local ordering will do as well: One could, for example, take the order from "first occurrence in the sentence under consideration." In this case, a would go with x_2 and b with x_1 . No matter: In any case, since what "satisfies" a sentence is a sequence of entities from the domain, one must have *some* ordering of the variables in order to be able to say which entity goes with which variable.

The first step in our extraction of the assignment parameters is to observe that, if we take what satisfies a sentence to be a sequence, there is *always* a detour through some sort of imposition of an order on the variables. This detour has as its only purpose defining which entity in the sequence goes with which variable. We avoid the detour by letting the second argument of the satisfaction relation be not a sequence but a *function*, a , defined on the variables, such that a *directly* assigns an individual in the domain to each variable, without presupposing a sequencing of any kind. So after this first step we have, for example,

$$Fx_2x_1 \text{ is satisfied by } a. \quad (2)$$

The second step is both linguistic and conceptual. Linguistically, we merely reword the very same relation. Instead of " a satisfies Fx_2x_1 ," we write

$$Fx_2x_1 \text{ is true at (or "on" or "with respect to" or whatever) } a. \quad (3)$$

The conceptual aspect of the second step is that in choosing (3) we are emphasizing that an expression such as " Fx_2x_1 " is grammatically sentential. We know it is sentential because it is subject to exactly the same embedding operations as any other sentence: conjunction, negation, and the like. Therefore, for example, " Fx_2x_1 " deserves a sentential semantics, which is to say, an account of the conditions under which it is *true*, in this case not absolutely, but relative to

an assignment of values to the variables. The language of satisfaction obscures this.

We call *anything* to which truth is relativized a *parameter* of truth (§6B.1). Once we see the need to relativize truth to assignment parameters, it is natural to consider other parameters. What in addition do we need to fix a truth value for, say, $\exists x_2 Fx_1x_2$? Evidently we need both an interpretation of F specifying the ordered pairs to which it truly applies, and also a domain specifying the range of the quantifier, $\exists x_2$. These are additional parameters of truth. The assignment parameters are, however, quite different from the interpretation and domain parameters. What makes them different is that quantifiers such as $\exists x_2$ are *translocal* in the assignment-to- x_2 parameter in the sense that in order to fix the truth value of $\exists x_2 A$ relative to a particular assignment to x_2 , one must sometimes look at the truth value of the embedded sentence, A , relative to assignments *other than* the assignment with which one starts. The assignment-to- x_2 parameter is therefore *mobile* in the sense that the language contains operations, namely, quantifiers such as $\exists x_2$, that are translocal in the assignment-to- x_2 parameter. In contrast to the quantifiers, the truth functions are one and all *local* in each assignment parameter: To fix the value of, for example, $\sim A$ at a family, a , of assignments to variables, one needs only the value of the embedded sentence, A , at a itself. And when one thinks of the domain and the interpretation of constants as parameters of truth, then these parameters are *immobile* if the language, as is typical or even universal, contains no operation that is translocal in them.

The key concepts of the preceding paragraph will be useful in speaking clearly about the language of indeterminism. "Mobile parameter" and "immobile parameter" were previously defined; see Definition 6-1 on p. 143. Here we insert definitions of the relations "local in" and "translocal in" between connectives and parameters. In order to avoid excess detail, however, we offer only approximations to the exact concepts.

8-1 DEFINITION. (*Local/translocal and mobile/immobile*)

- A connective (or any grammatical operation), Φ , is *local* in a parameter if in fixing whether $\Phi(A)$ is true or false at a certain value of that parameter, one needs to consider the truth or falsity of A at *only* that value.

Examples. $\forall x_1$ is local in the assignment-to- x_2 parameter. Truth functions are local in every parameter. Every operation is local in the *Domain* parameter.

- A connective, Φ , is *translocal* in a parameter if in fixing whether $\Phi(A)$ is true or false at a certain value of that parameter, one must consider the truth or falsity of A at *other* values of the parameter.

Example. $\forall x_1$ is paradigmatically translocal in the assignment-to- x_1 parameter.

Before moving on to linear time, branching time, and indexicals, let us mention a parameterization in our neighborhood that we do not discuss. *Modal logic*

characteristically adds one or two new immobile parameters, the set of “worlds,” and perhaps a “relative possibility” relation on that set, and also a mobile parameter, namely, the “world” parameter. The familiar modal connectives are translocal in the world parameter. This relativization of truth to worlds is not needed for understanding indeterminism as we conceive it, and so we add no such parameters. On the other hand, the technical devices and conceptual ideas of modal logic due especially to Kripke are essential to both plain or linear tense logic, and to the historical modalities of branching time.

8A.2 Prior’s linear tense semantics

Just as Tarski taught us how to interpret quantifiers by adding special parameters of truth (§8A.1), Prior has in the same way shown us how to understand much about tenses. *Linear tense logic* (generally called simply “tense logic”) is due to Prior 1957. On the semantic side, linear tense logic characteristically adds three new parameters to which to relativize truth: an immobile set of “times,” an immobile linear “temporal” ordering of this set, and a mobile “time of evaluation” parameter. All operations are local in the set of times and the temporal ordering, whereas the familiar tense connectives such as “it was true that,” “it will be true that,” and “it has always been true that” are translocal in the time parameter. We do not much discuss linear tense logic as a separate topic, since its ideas appear willy-nilly in the semantics for branching time. Reason: Each history is linearly ordered, so that we will see its workings in that context. For more information about linear tense logic, see Burgess 1984.

8A.3 Prior-Thomason’s branching-time semantics

We have seen that we learn from Tarski how to interpret quantified sentences (§8A.1), and from Prior how to understand linear tenses (§8A.2), in each case by thinking of truth as relativized to certain parameters. In exactly the same way, we learn from Prior 1967 and Thomason 1970 that if one wishes to be coherent, one has to interpret the subtleties of English tenses with the help of yet additional parameters of truth based on branching time. No philosopher who wishes to argue either for or against branching time should remain ignorant of this work; the penalty is contamination of an otherwise responsible appraisal by an almost certain—but unnecessary—tendency to slip into tensed modal muddles. We spell out some of the details of the Prior-Thomason semantics.

The indispensable idea from Prior-Thomason is that *truth shall be relativized to moment-history pairs*, where the moment belongs to the history. So, in addition to the immobile tree structure itself, there are two new mobile parameters to which truth (as well as denotation, etc.) is relativized: *the moment of evaluation*, and *the history of evaluation*. Linear tense operations are translocal in the former, shifting from moment to moment; and the historical modalities are translocal in the latter.

The most difficult part of this Prior-Thomason idea, the part that goes beyond linear tense logic, is negative in character:

Given indeterminism, it does not suffice to think of truth (or denotation, etc.) as relative only to moments.

One may wish that this were not so, but (in our opinion) it is. One must relativize truth to the *history* parameter as well. The reason is that only thus can we make sense, in branching time, out of plain (linear) future-tense sentences such as

There will be a sea battle tomorrow. (4)

Think of (4) as uttered before the admirals have made their decisions. Then the truth of that sentence (given indeterminism) depends not only on the moment at which the sentence is uttered. It depends in addition on which future course of events—which history—is being considered. (In §6D we investigated a variety of futile shortcuts.) To put the matter in easily understood words: Given indeterminism, what will happen (for example, whether or not there will be a sea battle tomorrow) depends on what will happen (i.e., on which history is being considered).

In yet other words, a sentence such as (4) depends on the history-of-evaluation parameter in exactly the same sense that a sentence like

x_1 is an admiral (5)

depends on the assignment-to- x_1 parameter. In each case, one has an example of “dependence” on a parameter in the sense of Definition 6-4, p. 153.

8A.4 Kaplan’s indexical semantics

Kaplan 1989 explains with maximum lucidity how indexical expressions should be semantically understood. (All page references are to this work by Kaplan.) The indispensable idea extracted from Tarski is that truth shall be relativized to assignments of values to the variables (§8A.1). The indispensable idea from Prior’s linear tense logic is that truth shall be relativized to time (§8A.2). And as we have just seen, the indispensable idea from Prior-Thomason is that truth shall be relativized to moments and histories (§8A.3). In the same way, the indispensable idea from Kaplan is that *truth shall be relativized to a context of use*.¹ Kaplan uses the context of use to determine much: the speaker of the context (used to obtain a denotation for “I”), the place of the context (used for “here”), the “world” of the context (used for “actually”), and the time of the context (used for “now”). These we can conveniently label as “*context parameters*.” They are brand new, and not to be identified with the mobile history or moment parameters. The latter were introduced to help with some translocal connectives, but the context parameters are used for indexicals.

To avoid complications irrelevant to indeterminism, we omit consideration of speaker and place. Of more significance is that we replace world and time with

¹Warning: Actual utterance is not what is at stake. For example, we wish to evaluate all the sentences of an argument in the same context of use. Kaplan explains this; see p. 546.

a single context parameter, *the moment of use*. If someone utters something, then (ideally) there is a unique moment to which that utterance is tied. Each such utterance has a unique causal past, and a unique future of possibilities, the whole of which is summed up by the moment of use. Among indexicals to which the moment of use contributes are *Now*: and *Actually*:

So the assignment parameters come from Tarski, the moment-of-evaluation and history-of-evaluation parameters from Prior-Thomason, and the moment-of-use parameter from Kaplan. We need two more families of parameters: the “structure parameters” and the “interpretation parameters.” Were our semantic purpose to stop with “truth,” although we would in that case certainly require these parameters, we could leave them implicit, as does Tarski, and indeed as we do in many places in this book. Leaving them implicit in the recursive account of “truth” would be warranted by the observation that the structure and interpretation parameters are immobile, so that they never need to be varied in the course of the recursion explaining truth. “Truth,” however, is not our only goal. In order to introduce ideas of “semantic equivalence” and “semantic implication” and “semantic validity,” we shall need to quantify over values of the structure and interpretation parameters. For this further purpose, we use the next two brief sections to be explicit about them.

8B Structure parameters: The “world” of the speakers

In order to fix truth, one needs some “*structure parameters*” to represent various features of the “world” of the speakers. We shall need parameters for the set of moments, for the ordering of these moments, for the set of instants, for the set of agents, for the choice function, and for the domain of quantification. When we gather these individual parameters, we call the result “*the structure parameter*.” A *structure* is a value of the structure parameter. A typical structure will look like $\langle \textit{Tree}, \leq, \textit{Instant}, \textit{Agent}, \textit{Choice}, \textit{Domain} \rangle$, and we let \mathfrak{S} range over such structures.

The key definition is this: $\mathfrak{S} = \langle \textit{Tree}, \leq, \textit{Instant}, \textit{Agent}, \textit{Choice}, \textit{Domain} \rangle$ is a *BT+I+AC structure* iff \mathfrak{S} satisfies the postulates on branching time with agents and choices summarily listed in §3, postulates that it was the business of chapter 7 to discuss in detail. “*BT+I+AC*” stands for “branching time + instants + agents and choices,” or, as we chiefly say, “agents and choices in branching time with instants” (leaving it to context whether or not *Domain*, needed only for quantification, is included).

In other places in this book we also discuss other structures of various kinds. The chief cases are these: (i) a *BT structure* (“*BT*” for “branching time”) has the form $\langle \textit{Tree}, \leq, \textit{Domain} \rangle$, (ii) a *BT+I structure* has the form $\langle \textit{Tree}, \leq, \textit{Instant}, \textit{Domain} \rangle$, and (iii) a *BT+AC structure* has the form $\langle \textit{Tree}, \leq, \textit{Agent}, \textit{Choice}, \textit{Domain} \rangle$. See §2 for a list of all of the kinds of structure that we treat.

In each case we implicitly allow that *Domain* might or might not be present, depending on whether or not quantification is at issue.

8C Interpretation and model: The “language” of the speakers

Parallel observations hold for the “language” of the speakers: What has elsewhere been implicit must here be made explicit. For this reason, we introduce “*the interpretation parameters*,” the value of each of which fixes an appropriate “meaning” for a single “nonlogical” constant. We assemble values of all the interpretation parameters into a single value, called an *interpretation*, of “*the interpretation parameter*.” We shall let \mathcal{I} be an interpretation. \mathcal{I} must represent the nonlogical atomic features of the “language” of the speakers by fixing a meaning of an appropriate type for each atomic nonlogical constant.² We say what we mean by “appropriate type” only for $BT+I+AC$ structures, leaving other cases to trivial adaptation.

In order to carry out this task, we first offer a preliminary account of the *grammar* of the language(s) of indeterminism that we propose to discuss. This is intended to be the same mini-language that we introduced as **L** in §6B.1 (with the *caveat* of note 4 on p. 141 that our account will be loose). In this chapter, however, we hardly ever use the denomination **L**. We call on some atomic features, and some modes of composition. Here we indicate only atomic features, composition by operators, and predication, as given in slightly more detail in §8. We leave the reader to infer the sentential modes of composition (the connectives) from the semantic clauses that we give in §8F. Typically any one of our discussions draws on only part of the following equipment: propositional variables p ; individual constants u (some of which are special terms for agents α , and one of which, \dagger , denotes “the non-existing object” to be available as a throw-away value of definite descriptions); individual variables x_j ; operator letters f ; and predicate letters F .

INTERPRETATION. (*Definition. Reference: Def. 15*) \mathcal{I} is an \mathfrak{S} -*interpretation* for a $BT+I+AC$ structure $\mathfrak{S} = \langle Tree, \leq, Instant, Agent, Choice, Domain \rangle$ iff \mathcal{I} is a function defined on propositional variables p , individual constants u (some of which are agent terms α and one of which, \dagger , is to denote “the non-existing object”), operator letters f , and predicate letters F , such that \mathcal{I} assigns to each propositional variable a function from *Moment-History* into $\{\mathbf{T}, \mathbf{F}\}$; assigns to each individual constant (including \dagger) a member of *Domain*; assigns to each agent term a member of *Agent*; assigns to each n -ary operator letter a function

²We tend to think of the “atomic” features of the language, except for individual variables, as “nonlogical,” and the compositional features as “logical.” In this study, however, we use “logical” and “nonlogical” without analysis and without serious commitment. We do refer the reader to MacFarlane 2000 for an argument that important features of the language of something like branching time should perhaps be classified as “logical” due to their proposed connection with the concept of assertion.

from *Moment-History* into functions from $Domain^n$ into $Domain$; and assigns to each n -ary predicate letter a function from *Moment-History* into functions from $Domain^n$ into $\{\mathbf{T}, \mathbf{F}\}$.

Representation of "world" and "language" must fit, and when they do, we call the combined representation a "model."

MODEL. (*Definition. Reference: Def. 15*) \mathfrak{M} is a *BT+I+AC model* (based) on \mathfrak{S} iff \mathfrak{M} is a pair $\langle \mathfrak{S}, \mathfrak{J} \rangle$, where \mathfrak{S} is a *BT+I+AC structure*, and where \mathfrak{J} is an \mathfrak{S} -interpretation.

We repeat that every one of the structure and interpretation parameters is immobile (the language has no operations translocal in any of the structure or interpretation parameters). We also repeat that their immobility is a principal reason that elsewhere we have been able to assume that the representations \mathfrak{S} of features of "world" and \mathfrak{J} of features of "language," encoded in the idea \mathfrak{M} of a model, are fixed implicitly in one particular way that makes it plausible that they are adequate idealized representations of certain features of *our* world and *our* language.

Grouping parameters of truth. We now have all the parameters we need. We rely on Definition 6-1 in order to characterize them as immobile or mobile. As indicated there, the structure and interpretation and moment-of-context parameters are immobile, whereas the assignment, moment-of-evaluation, and history-of-evaluation parameters are mobile.

8D Points of evaluation, and policies

Truth is relative to a specification of each parameter. Once we have fixed the values of the structure and interpretation parameters by specifying a model, $\mathfrak{M} = \langle \mathfrak{S}, \mathfrak{J} \rangle$, we still need to fix a number of other parameters: all the assignment-to-variable parameters, the moment-of-use parameter, the moment-of-evaluation parameter, and the history-of-evaluation parameter. Not just any fixing will do, as we indicate in the following definition of a "point."

POINT. (*Definition. Reference: Def. 16*) A *BT+I+AC point* is a tuple $\langle \mathfrak{M}, m_c, a, m/h \rangle$, such that \mathfrak{M} is a *BT+I+AC model*, $m_c \in Tree$, a is a function from the individual variables into $Domain$, $m \in h$, and $h \in History$. Henceforth we assume that $\langle \mathfrak{M}, m_c, a, m/h \rangle$ is a *BT+I+AC point*. We speak of the various parameters in $\langle \mathfrak{M}, m_c, a, m/h \rangle$ in the following way: \mathfrak{M} is the *BT+I+AC model*, m_c is the *moment of use*, a is the *assignment* (of values to the variables), m is the *moment of evaluation*, and h is the *history of evaluation*.

In this usage, “point” is short for “point of evaluation.” A point encodes a specification of the structure and interpretation parameters (via \mathfrak{M}), of the context, and of each of the three mobile parameters of truth.³

Sometimes, in order to avoid too many symbols, we use the following convention that lets us use π for $\langle \mathfrak{M}, m_c, a, m/h \rangle$.

π FOR $\langle \mathfrak{M}, m_c, a, m/h \rangle$. (*Definition. Reference: Def. 16*) We let π be the point $\langle \mathfrak{M}, m_c, a, m/h \rangle$, and adopt the following convention: In any context in which we write “ π ,” we will understand the expressions *Tree*, \leq , *Instant*, *Agent*, *Choice*, *Domain*, \mathfrak{S} , \mathfrak{J} , \mathfrak{M} , m_c , a , m , and h just as if we had written “ $\langle \mathfrak{M}, m_c, a, m/h \rangle$.”

It is striking that each point contains two independent references to moments: the moment of use, m_c , and the moment of evaluation, m . They play of course entirely different roles: The moment of use is immobile (unshifted by any operation), and may be used to fix indexical expressions, whereas the moment of evaluation is paradigmatically mobile, being shifted by a variety of tense constructions.

In addition to their uses in connection with various special indexical expressions, some context parameters have another role to play (Kaplan, p. 595; see our overlapping discussion in §6B.4.2). In the present case the following, which restates Policy 6-2, is critical.

8-2 OBSERVATION. (*Starting evaluation of stand-alone sentences*)
The moment-of-use parameter is used to start the evaluation of any stand-alone sentence considered as being uttered to some purpose involving the semantics of the sentence, for example, uttered as an assertion.

We expand the discussion of §6B.4.2. If you want to evaluate

Themistocles was surprised, (6)

and you understand that “was” moves evaluation into the past, you need a place to start that motion. When (6) is itself being considered as stand-alone, the moment of use gives us that starting point. It works in a special way: The (paradigmatically mobile) moment of *evaluation* is fixed, to begin with, as the very moment of *use*. It is only when we come to the sentence

Themistocles is surprised, (7)

which is implicitly embedded in (6) by past-tensing, that there is a divergence between moment of *use* (which remains the same as for (6) taken as stand-alone) and moment of *evaluation* (which needs shifting, existentially, toward the past).

³Elsewhere in this book, in order to minimize complexity of exposition, we not only keep the model, \mathfrak{M} , implicit; we also pass over recognition of either the assignment-to-variable parameters or the context parameters. In short, in spite of thinking of just one target language, we relativize truth in different ways depending on which problems we are attacking. For example, when we are not concentrating on quantifiers, we omit relativization to the assignment parameters.

You will note that in expressing Observation 8-2, which is derived from Kaplan, we use loose language. It is difficult to do otherwise for the following reason: Although *we* have in mind that a sentence can be considered either as stand-alone or as embeddable, the symbolic language (unlike, e.g., English) makes no such distinction. In the symbolic language, there is no syntactic mark (such as initial capitalization and final period in written English, or intonation in the spoken language) that distinguishes sentences taken as stand-alone from those taken as embeddable. This lack of match between English and the symbolic language makes analysis more difficult. Here is the best we know how to do (without describing a new kind of symbolic language) by means of a definition and a policy. (See Green 1998 for a study of “illocutionary-force-indicating devices,” including Frege’s sign of assertion.) First the definition, which is a *BT + I + AC*-specific version of the more general concept indicated in Policy 6-2.

CONTEXT-INITIALIZED POINT. (*Definition. Reference: Def. 16*) We say that a *BT + I + AC* point, $\langle \mathfrak{M}, m_c, a, m/h \rangle$, is *context initialized* iff the moment of use, m_c , is identical to the moment of evaluation, m . That is, context-initialized *BT + I + AC* points have always the form $\langle \mathfrak{M}, m_c, a, m_c/h \rangle$.

The idea is that in a context-initialized *BT + I + AC* point, the mobile moment of evaluation is “initialized by” the moment of the context of use. Next the policy.

8-3 POLICY. (*Differential treatment of stand-alone and embeddable sentences*)

- We recommend and urge, on pain of confusion, that sentences considered as *stand-alone* may usefully have their evaluation restricted to context-initialized points.
- We recommend and urge, on pain of confusion, that sentences considered as *embeddable* shall not have their evaluation so restricted, but that they shall be evaluated also at points in which the moment of use, m_c , and the moment of evaluation, m , diverge.

The rationale for the first part of this policy is that each utterance should be conceived as tied to a concrete context, and that such a context determines a unique causal position, with a definite past and a definite future of possibilities. We idealize such a position with the moment of use, m_c . This moment of use is the very moment at which we wish to evaluate a stand-alone sentence. That is why, for stand-alone sentences, we initialize the moment of evaluation with the moment of use. Keep in mind, however, that the moment of evaluation is mobile, and can be shifted by tense constructions ingredient in the stand-alone sentence. And that is the very reason for the second part of Policy 8-3.

Since there is in the symbolic language (and indeed often in the English examples of philosophers) no syntactic difference between stand-alone and embeddable sentences, the definition of “context-initialized point” (**Def. 16**) and

Policy 8-3 represent the best that we can do. The definition and policy are, we think, exceptionally useful in discussing tenses and indeterminism, for in those ventures the failure to observe the distinction between stand-alone and embeddable sentences is especially harmful.

Here is one way the definition and policy offer immediate progress. Kaplan (pp. 505–506) asks that in thinking about the “character” or meaning of a sentence, we first fix *context*, and then ask for the *content* of the sentence in that context. There should, however, be *two* notions of content-in-a-context, depending on whether we are thinking of the sentence as stand-alone or embeddable. If we are thinking of it as stand-alone, then the moment of evaluation is *initialized by* the moment of the context. Since one cannot reasonably treat a sentence with free variables as stand-alone (Assertability thesis 6-7), it is obvious that for stand-alone sentences the history is the *only* mobile parameter that is left to vary when considering a stand-alone sentence. If one correlates *time* to *moment* and *world* to *history*, this explains the otherwise puzzling phenomenon noted by Kaplan on p. 546:

... the truth of a proposition is not usually thought of as dependent on time as well as a possible world. The time is thought of as fixed by the context.

That is right for stand-alone sentences: Time (or moment) is fixed, while world (or history) is not. If, however, we are thinking of the sentence as embeddable by means of translocal connectives such as tense operators, then for “content in context” we *must* let (i) the moment of evaluation diverge from and vary independently of (ii) the moment of the context. This explains why Kaplan permits content to vary over times as well as “worlds.” We shall remain unclear as to the point of our semantic constructions unless we bear this in mind.

We note that Kaplan is working in a $T \times W$ framework (see §7A.6). It is, we think, an indication of the relative helpfulness of the moment-history framework that it explains a phenomenon that from the point of view of $T \times W$ seems just puzzling.

8D.1 History of the context?

Policy 8-3 arises partly in virtue of the fact that there is a pairing of two parameters that have the same range of variation, namely, the moment of context and the moment of evaluation. This is sometimes called “double indexing.” The phenomenon in general is of no special interest; after all, each x_j parameter in quantification theory has exactly the same range as any other, so that, for example, x_1 and x_2 exhibit “double indexing” of the domain. When, however, as in the case of moments, one of the paired parameters is a context parameter and the other a mobile parameter, we may speak more particularly of “context-mobile pairing.” We deepen our appreciation for Policy 8-3 if we ask the following two “context-mobile pairing” questions, one about histories, and one about assignments to variables. (The following discussion expands on that of §6B.5.)

8-4 QUESTION. (*History of the context?*) Because we shall have “modal” connectives that are translocal with respect to histories, it follows that there must be a mobile “history of evaluation” parameter. Furthermore, the mobility of the history of evaluation plays an essential role in our account of assertion (Semantic account 6-12); but why is there not also a “history of the context,” to be paired with the mobile history of evaluation, as *context* parameter? If there were, for stand-alone sentences we could *initialize* the history of evaluation with “the history of the context.”

It would clearly make *technical* sense to provide a context parameter ranging over histories, to be interpreted as “the history of the context of use.” And if we had one, we could enlarge our definition of “context-initialized point” to recognize this pairing. Further, there is plausibility in the Kaplan intuition that with “a little ingenuity” one can always make sense out of pairing a context parameter with a content parameter (p. 511).

Indeterminism, however, compels a view absolutely contrary to this: There is no “history of the context.” When you utter something, you do not thereby uniquely determine the entire future course of history. Your utterance has many choices and chances ahead of it, and so belongs to many histories. The context of use determines a unique moment, but not a unique history.

Just to be explicit, we mean to be challenging principles such as that suggested by Kaplan on p. 597: “Any difference in world history, no matter how remote, requires a difference in context.” Turning this around says that the identity of the context of use is enough to fix the course of world history, both past and future. That, if indeterminism be true, holds well enough for *past* history: The past, though largely unknown, is fixed. But it fails for the future: A single, well-identified context of use is typically part of a large variety of possible future courses of history. There is no unique “future of the context.”

The note on pp. 334–335 of Salmon 1989 is similar to Kaplan, though more convoluted. Salmon specifies the “quasi-technical notion of the context of an utterance,” to be distinguished from the “utterance” itself, by saying that

if any facts had been different in any way, even if they are only facts entirely independent of and isolated from the utterance itself, then the context of the utterance would, ipso facto, be a different context,

even if the utterance itself remains exactly the same. Salmon concludes that

... although a single utterance occurs in indefinitely many different possible worlds, any particular possible context of an utterance occurs in one and only one possible world.

If, however, “facts” are what is fixed at the moment of utterance, whether they are “isolated” or not, they cannot fix what the future brings—if indeterminism be true. Distinguishing “utterance” from a quasi-technical notion of “context of utterance” cannot make it otherwise. From this perspective, it looks as if Salmon is wrong—if indeterminism be true. It is, however, a delicate matter to label the

difference between Kaplan-Salmon and ourselves as a difference of opinion. Here is a place in which there is a dramatic difference between the ideas of "world" and of "history." Kaplan and Salmon and, for example, Lewis seem to picture a "world" as something like a single space-time. We, in contrast, picture a world as something like an indeterministic tree. What is *common* to these two pictures is the use of "world" for a system of events connected externally by some sort of causal-type relation. It is for this reason that on *either* picture, it is plausible that a given utterance, or context of utterance, will determine its one and only "world." This seems close to common ground. If, however, histories can branch indeterministically in the way that we suggest, then a single utterance, together with all the most distant "facts," belongs to many histories, no one of which is specially determined as "actual" by the moment of utterance.

Of course one could define "context of utterance" as a pair consisting of the moment of utterance together with a particular future history, the "actual" future history, and Salmon's note seems to suggest that he does in fact rely on the notion of an "actual" future history. He seems to rely, as do others, on the view that among all possible futures, one is marked out as a "Thin Red Line" in exactly the sense that we decried in chapter 6.

8D.2 Assignment-to-variables of the context?

As we discussed in §6B.5, there is still another family of mobile parameters, the assignment parameters, for which we do not provide a matching context parameter. We expand on that discussion.

8-5 QUESTION. (*Assignment-to- x_1 of the context?*) Because we shall have "quantifier" connectives that are translocal with respect to the assignment-to- x_1 (for example) parameter, there must be an "assignment-to- x_1 of evaluation" to be something like a "content" parameter, in which the quantifier is translocal. But why is there not also an "assignment-to- x_1 of the context" to be a paired *context* parameter? If there were, we could *initialize* the mobile assignment-to- x_1 parameter with the context assignment-to- x_1 parameter for stand-alone sentences.

It would clearly make *technical* sense to provide a matching assignment-to- x_1 as a context parameter. In some passages, indeed, Kaplan suggests that image (see especially p. 592), although he does not go so far as to provide *both* an assignment-to- x_1 context parameter *and* an assignment-to- x_1 mobile parameter in which the quantifiers can be translocal: There is only one assignment-to- x_1 parameter, not two.

Here the explanation lies not in the nature of things, but rather in (presumably universal) linguistic practice. Assignments to x_1 are not anchored in the context in any serious way. As Kaplan clearly says (p. 593), there is no "fact-of-the-matter" that contextually determines a unique assignment to x_1 . So the symbolic language we are describing fails to provide an assignment-to- x_1 as a context parameter because there is nothing in our language (or in any language we know) to which such a technical device would correspond.

It is good to recognize, however, that we speakers of English (supplemented with variables) *could* adopt a convention that supplied each variable with a value in each context of use. We *could* then go on to insist that, for stand-alone sentences, the mobile assignment used for evaluation is initialized by the context assignment. Suppose, for example, we require that the context determine that x_1 and x_2 are both assigned zero. That seems as good a convention as any, since there is (as Kaplan says) no “fact of the matter.” Then if we took something like

$$x_1 = x_2 \tag{8}$$

as a stand-alone sentence, and accordingly used Policy 8-3 to restrict ourselves to context-initialized points of evaluation, we should find out that (8) is automatically true. Ugh. Having realized that we *could* adopt such a convention, we are glad that we speakers of English haven’t done so. And the symbolic language we describe is, in this respect, just like ours. *All* of its context parameters are “fact of the matter” parameters really and objectively determined by the context of use; none are subject to some doubtful convention manufactured by a logician.

8D.3 Points and policies summary

(i) The special nature of context-initialized points is recognized. (ii) Policy 8-3 is firmly in place for differential consideration of stand-alone and embeddable sentences. (iii) That a point of evaluation, $\langle \mathcal{M}, m_c, a, m/h \rangle$, contains a context parameter for the moment of use, but none for “history of the context of use” and none for the “assignment of the context of use,” is no accident or idle logician-imposed convention. (iv) Nor is the fact that there is no analogy to Policy 8-3 for histories or assignments.

8E Generic semantic ideas

In the following section we detail recursive semantic clauses for the various operators that we treat. Here we outline generic features of the semantic concepts that are defined by that recursion. Except for the “in-context” concepts, these are all standard.

We are supposing two kinds of categorematic expressions, terms and sentences. The semantic value of a term is always an entity in *Domain*, while the semantic value of a sentence is always one of the two truth values, **T** or **F**. Since we are going to include operations that are translocal in assignments, moments, and histories, we know that we shall *have* to relativize semantic value to these parameters. We further explicitly relativize semantic value to the moment of context, even though it is immobile, partly for explicit indexicality, partly to exploit the idea of context-initialized points and the attendant Policy 8-3, but fundamentally because we are thinking of evaluating terms and sentences in many different contexts. And we explicitly relativize semantic value to

the structure and interpretation parameters because we wish to have not only “truth,” but also “equivalence,” “implication,” and “validity.”

We want a way of referring to the semantic value of any categorematic expression, be it term or sentence. Once we have “semantic value” for sentences, we automatically have truth, falsity, and the idea so important for indeterminism, settled truth.

SEMANTIC VALUE, DENOTATION, AND TRUTH. (*Definition. Reference: Def. 16*)

- *Semantic value.* For any categorematic expression, E , be it term or sentence, $Val_{\mathfrak{M}, m_c, a, m/h}(E)$, is “the semantic value of E at the point $\langle \mathfrak{M}, m_c, a, m/h \rangle$.” $Val_{\mathfrak{M}, m_c, a, m/h}(E)$ is defined recursively by clauses given in §8F and §8G. Note that by the earlier clause of **Def. 16** that appeared on p. 229, we may write $Val_{\pi}(E)$ in place of $Val_{\mathfrak{M}, m_c, a, m/h}(E)$.
- *Denotation.* Where t is any term, $Val_{\mathfrak{M}, m_c, a, m/h}(t) \in Domain$. $Val_{\pi}(t)$, or $Val_{\mathfrak{M}, m_c, a, m/h}(t)$, is “the denotation of t at the point $\langle \mathfrak{M}, m_c, a, m/h \rangle$.” Also, as before, $Val_{\pi}(t)$ stands in for $Val_{\mathfrak{M}, m_c, a, m/h}(t)$.
- *Truth.* $Val_{\mathfrak{M}, m_c, a, m/h}(A)$ is “the truth value of A at $\langle \mathfrak{M}, m_c, a, m/h \rangle$.” Where A is any sentence, $Val_{\mathfrak{M}, m_c, a, m/h}(A) \in \{\mathbf{T}, \mathbf{F}\}$.

Alternate much-used notation for truth and falsity:

- $\mathfrak{M}, m_c, a, m/h \models A$ iff $Val_{\mathfrak{M}, m_c, a, m/h}(A) = \mathbf{T}$. Either is read “ A is true at point $\langle \mathfrak{M}, m_c, a, m/h \rangle$.”
- $\mathfrak{M}, m_c, a, m/h \not\models A$ iff $Val_{\mathfrak{M}, m_c, a, m/h}(A) = \mathbf{F}$. Either is read “ A is false at point $\langle \mathfrak{M}, m_c, a, m/h \rangle$.”

We sharply distinguish settled truth, which is not history dependent, from plain truth, which is.

SETTLED TRUTH. (*Definition. Reference: Def. 17*)

- A is *settled true* at a moment m with respect to \mathfrak{M}, m_c , and a iff $\mathfrak{M}, m_c, a, m/h \models A$ for all $h \in H_{(m)}$. We may drop h , writing $\mathfrak{M}, m_c, a, m \models A$.
- A is *settled true throughout* a set of moments, M , with respect to \mathfrak{M}, m_c , and a iff $\mathfrak{M}, m_c, a, m \models A$ for all $m \in M$. In addition to dropping h , we may replace m by M , writing $\mathfrak{M}, m_c, a, M \models A$.
- We most often write just $\mathfrak{M}, M \models A$ and $\mathfrak{M}, m \models A$ since we use these concepts most often when assignment and context are not relevant.

Note that in order to avoid confusion, we do *not* introduce special notions of “truth” or “settled truth” for stand-alone sentences. We have two kinds of points, but not two kinds of truth or settled truth. We do, however, as indicated next, define two notions each of equivalence, implication, and validity, to be

used depending on whether a sentence is considered as embedded or as stand-alone. That is, we want strong notions of semantic equivalence, implication and validity, notions that are suitable for application to embeddable sentences and terms, and we also want versions suitable for stand-alone sentences. We pick out the latter concepts with the help of the adjective "in-context," which always signals restriction to context-initialized points. We define the ideas only for $BT+I+AC$ structures and models, leaving other cases for adaptation.

EQUIVALENCE. (*Definition. Reference: Def. 18*)

- Expressions E_1 and E_2 , either both terms or both sentences, are (*semantically*) *equivalent* iff for all $BT+I+AC$ points $\langle \mathfrak{M}, m_c, a, m/h \rangle$, $Val_{\mathfrak{M}, m_c, a, m/h}(E_1) = Val_{\mathfrak{M}, m_c, a, m/h}(E_2)$ (very same semantic value at all $BT+I+AC$ points). We write $E_1 \cong E_2$.
- Expressions E_1 and E_2 , both terms or both sentences, are *in-context equivalent* iff for all context-initialized $BT+I+AC$ points, $\langle \mathfrak{M}, m_c, a, m_c/h \rangle$, $Val_{\mathfrak{M}, m_c, a, m_c/h}(E_1) = Val_{\mathfrak{M}, m_c, a, m_c/h}(E_2)$ (the very same value at all context-initialized points). We write $E_1 \cong^{in-ctx} E_2$.

Semantic equivalence warrants replacement in any grammatical position, no matter how deep the embedding, and is thus of great importance. In-context equivalence, on the other hand, is suitable only for sentences considered as stand-alone, or occasionally for terms when considering replacement of one term by another in such a sentence.

IMPLICATION. (*Definition. Reference: Def. 19*)

- A set Γ of sentences *implies* a sentence A iff for all $BT+I+AC$ points $\langle \mathfrak{M}, m_c, a, m/h \rangle$, if $\mathfrak{M}, m_c, a, m/h \models A_1$ for every member A_1 of Γ , then $\mathfrak{M}, m_c, a, m/h \models A$ (truth preservation at all points). We write $\Gamma \models A$.
- A set Γ of sentences *in-context implies* a sentence A iff for all context-initialized $BT+I+AC$ points $\langle \mathfrak{M}, m_c, a, m_c/h \rangle$, if $\mathfrak{M}, m_c, a, m_c/h \models A_1$ for every member A_1 of Γ , then $\mathfrak{M}, m_c, a, m_c/h \models A$ (truth preservation at all context-initialized points). We write $\Gamma \models^{in-ctx} A$.

The strong notion of implication is important partly because it is monotone (or antitone) for each primitive connective in the language. That makes "logic," for instance natural deduction, easier. In-context implication is suitable only for stand-alone sentences.

VALIDITY CONCEPTS. (*Definition. Reference: Def. 20*)

- For \mathfrak{M} a model, A is *valid* in \mathfrak{M} (or *\mathfrak{M} -valid*) iff $\mathfrak{M}, m_c, a, m \models A$ for every $m_c \in Tree$ and a over *Domain* and $m \in Tree$. We write $\mathfrak{M} \models A$.
- For \mathfrak{S} a structure, A is *valid* in \mathfrak{S} (or *\mathfrak{S} -valid*) iff $\mathfrak{M} \models A$ for every \mathfrak{S} -model \mathfrak{M} . We write $\mathfrak{S} \models A$. When \mathfrak{S} is understood, especially when \mathfrak{S} is taken as a representation of *Our World*, we say that A is *valid*.

- When \mathbf{K} is a class of structures such as those listed in §2, A is *valid in \mathbf{K}* (or *\mathbf{K} -valid*) iff $\mathfrak{S} \models A$ for every structure \mathfrak{S} in \mathbf{K} .
- Each of these validity concepts has also an *in-context* version defined by restricting points to *context-initialized* points. When symbols are wanted, we write " \models^{in-ctx} " instead of " \models ". Thus, A is (i) in-context \mathfrak{M} -valid iff $\mathfrak{M} \models^{in-ctx} A$, (ii) in-context \mathfrak{S} -valid iff $\mathfrak{S} \models^{in-ctx} A$, and (iii) in-context \mathbf{K} -valid iff $\mathbf{K} \models^{in-ctx} A$.

Strong validity is interesting and useful, but perhaps not as intuitive as its in-context cousin. In-context validity is suitable only for sentences considered as stand-alone.

Observations on in-context semantic concepts. It is tempting to think that an in-context version of, for example, truth for stand-alone sentences is "real" truth. The temptation is to be firmly resisted: Truth is essentially parameterized, and in the end, one needs *all* the parameters. The reason is this: There is much insight to be gained by tracing the semantic value of expressions "from the inside to the outside," watching for the way that the semantic value of complex expressions depends on the value of their parts. For this enterprise, one should treat expressions as embeddable, and if so, "real" semantic value—including "real" truth—has to be relativized to the assemblage of all parameters. None can be left out for "real" truth.⁴

There is a difference between the strong semantic ideas and their in-context cousins for sentences that contain indexical expressions, such as *Now*., which pick up their meaning from the moment-of-use parameter. Here is an example, after Goodman, that relies on the difference between "strong absurdity" and its in-context cousin, "in-context absurdity." (We are thinking of absurdity in an intuitive way; our discussion does not require anything more.)

The boat Jack owns is larger than the boat Jack owns now. (9)

Although (9) is not strictly contradictory, it sounds absurd, and, as a stand-alone sentence, it *is* absurd. The (tedious) semantic explanation is this. There are two definite descriptions involved. For brevity, let B be the phrase, "the boat Jack owns," and let $Now-B$ be the phrase, "the boat Jack owns now." Semantically, the referent of B depends on the moment of evaluation, and is independent of the moment of utterance. The semantic properties of $Now-B$ are just the reverse: The referent of $Now-B$ depends on the moment of utterance (because of the "now"), but is independent of the moment of evaluation. In some sense, then, B and $Now-B$ might refer to different boats. (B is not semantically equivalent to $Now-B$.) But if we are thinking of (9) as stand-alone, then we are to invoke in-context concepts: In checking for absurdity, we are to evaluate (9) at only context-initialized points. At context-initialized points, the moment

⁴It is part of our meaning that one should not think of "supervaluation" as giving "real" truth. Here we depart from the recommendation of Thomason 1970.

of evaluation is identical to the moment of utterance, so that at all context-initialized points, B and $Now-B$ must refer to the same boat. (B is in-context equivalent to $Now-B$.) No boat can be larger than itself; whence the in-context absurdity.

This result is good, because when you are thinking of (9) as stand-alone, the refined notion of in-context absurdity seems in compelling correspondence with one's intuitive judgment of absurdity. If, however, you are thinking of (9) as embeddable, then you do not care that it is in-context absurd. You care only that it is *not* absurd in the strong sense that refers to *all* points, not just to the context-initialized points. If (9) were absurd in the strong sense, then the following, which embeds (9) within a future tense, would also be absurd.

It will be true that the boat Jack owns is larger than the boat
he owns now. (10)

Since (10) makes perfectly good sense, it must be that its contained part, (9), is not absurd. The (tedious) semantic explanation is this. The denotation of $B =$ "the boat Jack owns" depends on the moment of evaluation, and is independent of the moment of utterance. The semantic properties of $Now-B =$ "the boat Jack owns now" are just the reverse: The denotation of $Now-B$ depends on the moment of utterance (because of the "now"), but is independent of the moment of evaluation. The function of the future tense in (10) is precisely to *move* the moment of evaluation, while leaving the moment of utterance unchanged. ("Will" is translocal in the moment of evaluation and local in the moment of utterance.) You must evaluate (9) at points that are *not* context initialized: You must evaluate the embedded sentence, (9), with the moment of evaluation being future to the moment of utterance. And when you do that, B refers to the boat Jack owns at that future moment, while $Now-B$ still refers to the boat Jack owns at the moment of utterance. Since it makes perfect sense that these boats be distinct, it is not at all absurd to suppose that one is larger than the other. Hence, although (9) is absurd when taken as standing alone, when (9) is considered as embeddable, for example, in (10), it is by no means absurd, a fact that we explain by considering points that are *not* context initialized. For indexical-free sentences, however, there is no difference whatsoever between the strong concepts and the in-context concepts.

Suppose that A is being considered as stand-alone. Suppose that A is also without free variables, and is therefore assignment-independent. Then one could meaningfully omit both the moment of evaluation parameter *and* the assignment parameter, leaving only model, context, and history. One cannot, however, go on to omit the history parameter, even for stand-alone sentences. We *need* the relativization to histories if we are to understand a stand-alone sentence such as "There will be a sea battle tomorrow," which may well be asserted before the matter is settled true or settled false. See especially Assertability thesis 6-7.

8F Semantics for stit-free locutions

To complete the semantics for the language of indeterminism, we need to give explanations of the semantic values of atomic terms and sentences, and of how the semantic values of more complex expressions arise out of the semantic values of their “parts.” We follow Curry in calling a way of building new expressions out of old a “functor.” We treat functors that build sentences from sentences (connectives), that build sentences from terms (predicates), that build terms from terms (operators), that build terms from sentences (subnectors), and that build sentences from a combination of terms and sentences (mixed nectors). Context will always make it clear which sort of functor is at stake. Mathematically we are defining “ $Val_{\mathfrak{M}, m_c, a, m/h}(E)$ ” (the semantic value of E) by simultaneous recursion on terms and sentences. Conceptually it seems better (or anyhow at least as good) to take the following as explaining the meaning of the various features of the language of indeterminism in terms of some prior understanding of “ $Val_{\mathfrak{M}, m_c, a, m/h}(E)$.”

For brevity in stating upcoming semantic clauses involving variable-binding functors, we adopt the following.

8-6 DEFINITION. (*Semantic abbreviations*)

- “ x_j ” can name the assignment-to- x_j parameter. (We also continue to use “ x_j ” to name a piece of notation.)
- We want a short way of “shifting” the values of the assignment parameter. Let z be any appropriate value of the x_j parameter, which is to say, let $z \in Domain$; and let $\langle \mathfrak{M}, m_c, a, m/h \rangle$ be any point. Then

$$[z / x_j] \langle \mathfrak{M}, m_c, a, m/h \rangle$$

is the point that is just like $\langle \mathfrak{M}, m_c, a, m/h \rangle$ except that the value of the parameter, x_j , is shifted to be z . (If there is no such point, the notation is undefined.)

In the rest of this long section we present a plethora of semantically explanatory clauses for the “stit-free” part of the language. Then in §8G we go over the stit ideas. We organize the work of the present section on the “stit-free” part as follows. After the clauses covering the atomic expressions of the language, we organize our explanations in terms of abstract properties of the functors, primarily considering whether they are “local” or “translocal” in various mobile parameters in the sense of Definition 8-1. Also we occasionally need the concept of “anchoring.” Before proceeding, we give a definition of “anchoring” in rough terms, which suffices for present rough purposes. (The rigorous definition is too tedious to be helpful here.)

8-7 DEFINITION. (*Anchoring*) Let Φ be a functor.

- Φ is *anchored* in a parameter if in passing from E to $\Phi(E)$, you require the very identity of the current value of that parameter, over and beyond the pattern of semantic values of E as the parameter varies.

- Φ is *unanchored* in a parameter iff it is not anchored in it.

EXAMPLES. (*Anchoring*) Negation is not anchored in any parameter, since it cares only for the truth value of A , no matter the values of any parameter. The universal quantifier, $\forall x_j$, is unanchored in the x_j parameter (it needs only the pattern of truth values, not the actual value of x_j); but $\forall x_j$ is anchored in the *Domain* parameter (you need to know the actual domain). In modal logic, necessity is always anchored in the set-of-worlds parameter; in **S4** it is also anchored in the world parameter (you have to know where you are in order to quantify over the relatively possible worlds), whereas in **S5**, necessity need not be anchored in the world parameter (since in the case in which you are quantifying over all worlds, you do not need to know where you are). And, paradigmatically, indexicals are anchored in the context of use.

Having put the definitions of assignment-shifting and anchoring in play, we proceed to detailed semantic clauses, which will occupy us until the end of this chapter.

8F.1 Atomic terms and sentences; operators and predicates

We are considering a mini-language with the following equipment.

- p (and sometimes q) ranges over propositional variables.
- u ranges over individual constants, including two sorts of special terms: (i) α ranges over agent terms (and frequently over the agents themselves), and (ii) \dagger is a term that artificially denotes “the non-existing object,” to be available as a throwaway value of definite descriptions when existence or uniqueness fails.
- x_j ranges over individual variables.
- f ranges over operator letters.
- F ranges over predicate letters.

Terms and sentences.

- t ranges over terms of any kind. $f(t_1, \dots, t_n)$ is a term.
- A ranges over sentences. We also sometimes use B, C, D, P , and especially Q as ranging over sentences. $F(t_1, \dots, t_n)$ is a sentence.

The semantics of these features is as follows.

- For A an atomic sentence, $Val_{\mathfrak{M}, m_c, a, m/h}(A) = \mathcal{J}(A)(m/h)$.
- For t an individual constant u (including agent-terms α and the term, \dagger , for “the non-existing object”), $Val_{\mathfrak{M}, m_c, a, m/h}(t) = \mathcal{J}(t)$.

- For x_j a variable, $Val_{\mathfrak{M}, m_c, a, m/h}(x_j) = a(x_j)$.
- For n -ary operator f and terms t_1, \dots, t_n , $Val_{\pi}(f(t_1, \dots, t_n)) = \mathcal{I}(f)(m/h)(Val_{\pi}(t_1), \dots, Val_{\pi}(t_n))$.
- For n -ary predicate F and terms t_1, \dots, t_n , $Val_{\pi}(F(t_1, \dots, t_n)) = \mathcal{I}(F)(m/h)(Val_{\pi}(t_1), \dots, Val_{\pi}(t_n))$.

8F.2 Absolute functors

A functor is defined as “*absolute*” iff it is both local in and unanchored in every parameter. Identity, alone among predicates, has this property. Among connectives, only the truth functions are absolute.⁵ Here are the semantic explanations of a few absolute functors.

- $t_1 = t_2$. READING: t_1 is identical to t_2 . SEMANTICS: $\pi \models t_1 = t_2$ iff $Val_{\pi}(t_1) = Val_{\pi}(t_2)$.
- $\sim A$. READING: That A is not true. SEMANTICS: $\pi \models \sim A$ iff $\pi \not\models A$.
- $A_1 \& A_2$. READING: A_1 and A_2 . SEMANTICS: $\pi \models A_1 \& A_2$ iff $\pi \models A_1$ and $\pi \models A_2$.
- Other truth-functional connectives are analogous, or may be introduced with their usual abbreviations. We occasionally use the following: \vee , \supset , \equiv , \top and \perp .

8F.3 Variable-binding functors: Translocal in an assignment parameter

- $\forall x_j A$. READING: For all x_j , A . SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models \forall x_j A$ iff $\forall z[z \in \text{Domain} \rightarrow [z / x_j](\mathfrak{M}, m_c, a, m/h) \models A]$.
- $\iota x_j(A)$. READING: A definite description: the sole x_j such that A . SEMANTICS: $Val_{\mathfrak{M}, m_c, a, m/h}(\iota x_j(A)) =$ the sole z such that $[z / x_j](\mathfrak{M}, m_c, a, m/h) \models A$ if there is exactly one such. If not, $Val_{\mathfrak{M}, m_c, a, m/h}(\iota x_j(A)) = \mathcal{I}(\dagger)$ (“the non-existing object”).
- Other variable-binding operators are analogous.

⁵That makes it plausible to guess that “absolute” is a good thing to mean by “extensional.” That won’t work, however, because, e.g., predicate letters (of the sort considered here) are thought to carry only extensional meaning, even though they are anchored in the model, \mathfrak{M} , and therefore not absolute. One might next guess that “local in all parameters” is a good meaning for “extensional.” Observe, however, that if we were to take this as an explication of “extensional,” then quantifiers would turn out just as “non-extensional” as, e.g., the modalities. Given all this confusion, and since it is difficult to determine what “extensional” means (van Bentham 1988, p. 109, suggests that “no general satisfactory definition seems to exist”), we feel that it is better to stick to the ideas of “local” and “absolute,” which have clear and definite meanings that are *logical* rather than ideological.

Grammar $\iota x_j(A)$ is a (categorematic) term, the definite description. Its value when existence-and-uniqueness fails is a mere throwaway. It corresponds to Kaplan's definite description. For his "dthat" one may use $\iota x_j(\text{Actually}_1:A)$, which is context-dependent, and independent of the moment and history of evaluation. (We use ι instead of its inversion for convenience.)

8F.4 Historical-modality connectives: Translocal in the history parameter

These are the "historical-modality" connectives, closely tied to some of our intuitive ideas of necessity and possibility. In processing these connectives, you do not have to consider any other moment than the current moment of evaluation: They are local in the moment-of-evaluation parameter.

- *Sett:A*. READING: It is settled true, or historically necessary, that *A*. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models \text{Sett}:A$ iff $\forall h_1[(m \in h_1 \text{ and } h_1 \in \text{History}) \rightarrow \mathfrak{M}, m_c, a, m/h_1 \models A]$.
- *Poss:A*. READING: It is historically possible that *A*. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models \text{Poss}:A$ iff $\exists h_1[m \in h_1 \text{ and } h_1 \in \text{History} \text{ and } \mathfrak{M}, m_c, a, m/h_1 \models A]$.
- *Can:A*. READING: It can be that *A*. DEFINITION: *Can:A* iff *Poss:A*.

We use *Can:* as a mere stylistic variation of *Poss:*. This gives a version of the "all in" ability idea of Austin 1961 (p. 177) when combined with the deliberative stit as in §9G. For versions suitable for use with the achievement stit, see §8G.

8F.5 Tense and temporal connectives: Translocal in the moment-of-evaluation parameter

Next are "standard" linear temporal connectives of the sort introduced by Prior 1957. All of these connectives should be thought of as "temporal." *They are—this is fundamental—local in the history parameter.* That is, temporal connectives, including the tenses, always start on a given history, and the connective moves you, but only "vertically." *These connectives never move you off the history on which you started—that is what "local in the history parameter" means.* For this reason, they are entirely "linear," and have the same "logic" as standard tense logic. The use of these connectives is absolutely natural and familiar; they require no strange explanations—once one accepts that, in analogy with variables, there is always a current history of evaluation even though there is no history of the context. They interact, however, in enlightening ways with the historical modalities.

- *Was:A*. READING: It was true that *A*. Or put the main verb of the reading of *A*—if it has one—into the past tense. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models \text{Was}:A$ iff $\exists m_1[m_1 \in h \text{ and } m_1 < m \text{ and } \mathfrak{M}, m_c, a, m_1/h \models A]$.

Prior's "P." Shift along the present history, existentially, to earlier moments, and check them for A (with respect to the current history). Because histories are closed downward and thus form a chain, the clause " $m_1 \in h$ " is redundant. That the route of backward travel is uniquely determined, however, should not blind you to the importance of the nontriviality of keeping to the same (the current) history when you evaluate A . As Prior explained, $Was:A$ does not imply $Sett:Was:A$ —although $Was:Sett.A$ does indeed imply $Sett:Was:A$. The point is that although the route traveled is unique, it is part of many different histories

- $Will:A$. READING: It will be true that A ; or put the main verb of the reading of A , if it has one, into the future tense. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models Will:A$ iff $\exists m_1[m_1 \in h$ and $m < m_1$ and $\mathfrak{M}, m_c, a, m_1/h \models A$].

Prior's "F." Here you shoot forward, existentially, along the current history, checking each moment along the way. And look: In contrast to $Was:$, you cannot, in understanding $Will:$, get rid of a reference to the history parameter.

- $Was-always:A$ (Prior's H) and $Will-always:A$ (Prior's G) have analogous clauses that shift you, quantificationally, along the current history. $Some-time:A$ and $Always:A$ are similarly local in the history of evaluation.
- $At-inst_t:A$. READING: That A was, is or will be true at (the instant or time) t . When English A is not too complicated, " A at t " or " $At\ t, A$," will often do. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models At-inst_t:A$ iff $(Val_\pi(t) \in Instant$ and $\langle \mathfrak{M}, m_c, a, m_{(Val_\pi(t), h)} / h \rangle \models A)$.

Grammar: t is any term, and A is any sentence. $At-inst_t:A$ reflects English constructions such as "At 4:00 P.M. the coin will come up heads." The clause tells you to shift $\mathfrak{M}, m_c, a, m/h$ by replacing the current value of the moment-of-evaluation parameter by a new moment, namely, the one (and the only one) in which the current history intersects the instant specified by the value of t . "Travel up or down the current history until you hit the instant t ; that moment is where you must evaluate A (with respect to the current history)." This understanding of $At-inst_t:A$ makes it false when t does not refer to an instant. We manage to live with this awkwardness. You can see that $At-inst_t:$ is translocal in and only in the moment-of-evaluation parameter, for that is the only parameter that is shifted. Also, as long as t itself is independent of the moment-of-evaluation parameter, $At-inst_t:A$ is bound to be independent of the moment-of-evaluation parameter. For instance, let t be "4:00 P.M.," with some definite date understood. Adding additional temporal connectives to $At-inst_{4:00\text{P.M.}}:A$ has no more effect than (is just as vacuous as) nesting one x_j -binding operator within another. For example, "It has always been true that at 4:00 P.M. the coin comes up heads" just reduces to "At 4:00 P.M. the coin comes up heads." So does "At 2:00 P.M. at 4:00 P.M. the coin comes up heads" (see the discussion of Figure 6.3).

Some philosophical logicians feel that including terms naming times (or instants) in a formal tensed language is wrong-headed, or at best inelegant. Our excuse for doing so is that such terms as used in "at" constructions play a role in many philosophical discussions of determinism versus indeterminism. Not everyone is clear on how they should work under indeterminism, which makes it worthwhile to clarify their logic. Of particular importance in understanding indeterminism is the fact that $At-inst_{4:00\text{P.M.}}:A$ is not in general independent of the history-of-evaluation parameter. "At 4:00 P.M. the coin comes up heads," or even "It has always been true that at 4:00 P.M. the coin comes up heads," can be just as dependent on the history-of-evaluation parameter as "It will be true that the coin comes up heads." "At" constructions pose quite the same problems as does the future tense. These matters are difficult to hear or see in natural language,

where we can (apparently) say, "At 4:00 P.M. it will be inevitable that Jack at 1:00 P.M. might be running at 3:00 P.M. at 5:00 P.M." We don't know if this makes sense or not. Our point is that those who make philosophical points explicit in a Prior-Thomason connective form have much less difficulty—especially with regard to scope, which is so ambiguous in English and so clear when all tense talk is carried by connectives.

- $At\text{-}mom_t:A$. READING: That A was, is, or will be true at the moment t . SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models At\text{-}mom_t:A$ iff $Val_\pi(t) \in h$ and $\langle \mathfrak{M}, m_c, a, Val_\pi(t)/h \rangle \models A$.

Grammar: t is any term and A is any sentence. This construction travels up or down the current history to the moment denoted by t , and evaluates A just there, still on the current history. It does not move off the current history (it is local in the history parameter), and is therefore a true temporal construction. Accordingly, it cannot be successfully used in connection with moments off the current history. (Contrast: $At\text{-}inst_t:A$ is always successful when t denotes an instant, since the current history intersects every instant.) The $At\text{-}mom_t$: connective seems much stranger than the $At\text{-}inst_t$: connective, chiefly, we suppose, because we don't have names for moments, but we do have names for instants (or times). The $At\text{-}mom_t$: construction is nevertheless of great usefulness in untangling "double time references" such as are needed in order accurately to understand the following scenario. (Our account may sound a little awkward, because we omit several indexicals that would naturally be used.)

It's 4:00 P.M. At 2:00, Themistocles said, "I promise that Themistocles will choose to fight a sea battle." (11)

So in direct speech, what Themistocles promised was this:

Themistocles will choose to fight a sea battle. (12)

We are interested in how one can use the semantic content of (12) in order to illuminate promise keeping. So let it be true at 4:00 that Themistocles has kept his promise; such is, at 4:00, a settled fact. What does this mean? Here are two candidates that don't work.

- At 4:00 Themistocles has kept his promise iff (12) is settled true now, at 4:00. This is wrong. After all, now, at 4:00, what Themistocles promised, namely, (12), is settled false. The point is that when evaluated at 4:00, the future tense of (12) would reach forward into times later than 4:00, long after the sea battle. (We neglect as a distraction the possibility that there be *another* sea battle.)
- At 4:00 Themistocles has kept his promise iff (12) is settled true at the moment of promising. This is wrong. At the moment of promising, the sentence, (12), was *not* settled true. That is, (12) was, at the moment of promising, true on some histories and not on others—for the choosing still lay in the future.

What signifies that the promise has been kept is more complicated.

- At 4:00 Themistocles has kept his promise iff at 4:00 it is settled true that (12) was true *at the moment of promising*. The key is the "double time reference": The "settled" is evaluated later, at 4:00, while what is settled is the truth of (12) at the moment of promising. See §5C.2.1 for a more extended discussion of double time references.

The upshot is that you cannot do without “the moment of promising,” even though you certainly do not have a name for it. Here is how the analysis could go in indirect speech, where, as advertised, we can use the *At-mom_t* connective. (Also we shall be using (12) instead of mentioning it.) We let the moment of promising be m_p . At 4.00 Themistocles kept his promise iff

- Wrong: *At-inst_{4 00}*:(12).
- Wrong: *At-mom_{m_p}*;*Sett*:(12).
- Right: *At-inst_{4 00}*.*Sett*:*At-mom_{m_p}*:(12)

None of these connectives can be shifted around; and it is essential that there be a “double time reference.”

8F.6 Some mixed temporal-modal connectives

We add just a few “mixed” connectives. All these are definable, and in some cases we give a definition instead of a semantics.

- *Inevitably:A*. READING: It is inevitable (sooner or later) that *A*. DEFINITION: *Sett:Will:A*.
- *Was-always-inevitable:A*. READING: It has always been inevitable that *A* would be true at the present time. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models$ *Was-always-inevitable:A* iff $\mathfrak{M}, m_c, a, m_1/h_1 \models A$ for every member m_1 of $i_{(m)}$ and every h_1 to which m_1 belongs.

Note that *Was-always-inevitable:A* is not equivalent to *Was-always-Inevitably:A*, since the former makes reference to the present instant. It is useful in connection with the achievement stit.

- *Might-have-been:A*. READING: That *A* might have been true at the present time. DEFINITION: \sim *Was-always-inevitable:* $\sim A$.
- *Universally:A*. READING: It is universally true that *A*. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models$ *Universally:A* iff $\mathfrak{M}, m_c, a, m_1/h_1 \models A$ for every member m_1 of *Tree* and every h_1 to which m_1 belongs.

This connective crops up only infrequently in this book. Accordingly, so as to be maximally mnemonic, we refrain from abbreviating “universally.” Note that because of historical connection, *Universally:A* is equivalent to the following mouthful: *Was-always:Sett:Will-always:A*. When *A* has a syntax appropriate for a “lawlike statement,” and perhaps has also an appropriate social standing in the scientific community, the truth of *Universally:A* might be a sensible thing to mean by saying that *A* is a “law.” If there are other worlds, perhaps one could term such a law “a law of our world,” and imagine other worlds with other laws. In this sense, *Universally:A* could well be an “imaginative contingency,” an intentional idea for which we propose no theory.

8F.7 Indexical connectives: Anchored in the context of use

Both the “now” connective and the “actuality” connectives are indexical in the sense that they are anchored in a context-of-use parameter. (These connectives are also translocal in the moment-of-evaluation parameter, but that feature is not shared by all indexical constructions.)

- *Now:A*. READING: It is now true that *A*. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models \text{Now:A}$ iff $\mathfrak{M}, m_c, a, m_{(v(m_c), h)}/h \models A$.

Our indeterminist adaptation of Kamp’s well-known explanation of *Now:A* comes to this: Stay on the current history, but travel up or down to where the current history intersects the very instant (time) of the context of use. Using that context-determined moment with the current history, check the truth value of *A*. The new moment will be “at the same time as” the moment of use, but since it lies on the current history rather than on any history containing the moment of use, it may well be inconsistent with the moment of use, as in the following: “I am not now rich, but if the coin had come up heads, now I would be rich.” In symbols, with A_1 for “rich” and A_2 for “heads”: $\text{Now:}\sim A_1 \ \& \ \text{Was:}(\text{Poss:}A_2 \ \& \ \text{Sett:}(A_2 \supset \text{Now:A}_1))$.

- Actuality connectives. Actuality is hard, partly because different thinkers have such different views. We are getting at “actuality” in an indexical sense, shared by, for example, Kaplan and Lewis, that anchors actuality in the context of use. Of this species, the easiest-to-understand actuality connective, *Actually*₁; ties actuality to what is settled true at the moment of use. There is, however, an interesting variant, *Actually*₂; that has a disjunctive nature. On this variant, one considers whether or not the moment of use sits on the current history. If so, the variant ties actuality to what is *plain* true at the moment of use on the current history. If, however, the moment of use does not sit on the current history, the variant ties actuality to what is *settled* true at the moment of use.⁶ We mark the easier-to-understand and the disjunctive variant with a simple “1” and “2” respectively because we have little to say about their different properties and applications.

– *Actually*₁:*A*. READING: It is actually true that *A*. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models \text{Actually}_1:A$ iff $\forall h_1[m_c \in h_1 \ \& \ h_1 \in \text{History} \rightarrow \mathfrak{M}, m_c, a, m_c/h_1 \models A]$.

– *Actually*₂:*A*. READING: It is actually true that *A*. SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models \text{Actually}_2:A$ iff either $(m_c \in h \ \& \ \mathfrak{M}, m_c, a, m_c/h \models A)$ or $(m_c \notin h \ \& \ \forall h_1[m_c \in h_1 \ \& \ h_1 \in \text{History} \rightarrow \mathfrak{M}, m_c, a, m_c/h_1 \models A])$.

⁶It would be in the spirit of Thomason and Gupta 1980 to add a parameter that makes sense out of picking a *particular* history through the moment of use, perhaps a history that is “very like” the current history, which passes through the current moment. Since, however, “very like” is not an idea that fits well with the more austere causal notions that we employ, we forgo adding such a parameter.

As always, *Now*: and the actuality connectives do very little work at the head of sentences considered as stand-alone, since anyhow, by Policy 8-3, the moment of evaluation is already initialized by the moment of use, so that the shifting called for by *Now*: or an actuality connective is vacuous. When embedded in translocative connectives, however, *Now*: and the actuality connectives really do shift the moment of evaluation.

This completes our discussion of how various variable-binding, temporal, history-modal, and indexical functors work in branching time. Next we take up the stit functors.

8G Clauses for stit functors

We give properly recursive forms of semantic explanations of the choice-based concepts—the stit concepts—that have been previously defined and used in various places in this book.

8G.1 Semantics for the deliberative stit

The simplest of our two major stit concepts is the “deliberative” stit, as defined first by von Kutschera 1986 and later by Horty 1989, and which we apply at length in chapter 11 and chapter 12. Recall **Def. 16**, according to which $\pi = \langle \mathfrak{M}, m_c, a, m/h \rangle$.

- [*t dstit*: *A*]. **READING**: *t* sees to it that *A*. **SEMANTICS**: $\mathfrak{M}, m_c, a, m/h \models [t \text{ dstit}: A]$ iff the following conditions are satisfied.
 - *Agency*. $Val_\pi(t) \in Agent$.
 - *Positive condition*. For every h_1 , if $h_1 \in Choice_m^{Val_\pi(t)}(h)$, then $\mathfrak{M}, m_c, a, m/h_1 \models A$.
 - *Negative condition*. For some $h_2 \in H_{(m)}$, $\mathfrak{M}, m_c, a, m/h_2 \not\models A$. The history h_2 on which *A* is not true is sometimes called a “counter.”

The agency condition is needed when the agent position of a stit sentence is available to every term. In this book we do not often enter the agency condition explicitly, since usually, such as when we reserve “ α ” for an agent-term, it is a *presupposition* that the term in the agent position denotes an agent. We remark that in normal applications, one would expect that the term, *t*, is independent of both history and moment, as is always true for our use of α . One will usually obtain what one wants by trading in [*t dstit*: *A*] for $\exists x_1[x_1 = t \ \& \ [x_1 \text{ stit}: A]]$.

In the semantics of *dstit*, there is no double temporal reference: The moment of evaluation of *A* and the moment of choice are identified, so that while *dstit* is an entirely viable candidate for helping to analyze agentive locutions, its expected properties are considerably different from those of the achievement stit. If one wishes to think of *dstit* as reporting an “action” or even a “choice,” it is difficult to say in comfortable English just when the choice is made. The difficulty is that before or at *m* it is not yet settled which choice $Val_\pi(t)$ makes, while at any later moment the choice has already been made. The source of the difficulty is that *dstit* reports an *immediate* transition, **Def. 8**, so that there is no room for an “action” qua “event” between initial and outcome.

The solution is to understand “the action” as being the transition itself. Since a transition consists in a pair of “events,” initial event and outcome event, it is obvious that a transition cannot have a “simple location” (Whitehead).

8G.2 The Chellas stit

An even simpler stit concept, from Chellas 1992, omits the negative condition:

- $\mathfrak{M}, m_c, a, m/h \models [\alpha \text{ cstit}: A]$ iff $\mathfrak{M}, m_c, a, m/h_2 \models A$ for all h_2 with $h_2 \equiv_m^\alpha h$.

8G.3 Achievement-stit semantics based on witness by moments

In Definition 2-4, we presented a semantic explanation of the achievement stit. In this section we present those same semantics in just slightly different words. Then in §8G.4 we generalize in order to minister to a specific shortcoming, and in §8G.5 we suggest a potentially helpful alternative.

- $[t \text{ stit}: A]$. **READING:** t sees to it that A (the present fact that A is entirely due to a prior choice of t). **SEMANTICS:** $\mathfrak{M}, m_c, a, m/h \models [t \text{ stit}: A]$ iff there is a moment, w , that is a *momentary witness to* $[t \text{ stit}: A]$ at m relative to $\langle \mathfrak{M}, m_c, a, m/h \rangle$, where that phrase is defined by the following conditions.
 - *Agency.* $Val_\pi(t) \in \text{Agent}$.
 - *Priority.* w is properly earlier than m : $w < m$.
 - *Positive condition.* $\mathfrak{M}, m_c, a, m_1/h_1 \models A$ for every moment m_1 in $\text{Choice}_w^{Val_\pi(t)}(m)$, and for every history h_1 in $H(m_1)$.
 - *Negative condition.* w must lie under some moment m_1 in $i(m)$ such that there is a $h_1 \in H(m_1)$ such that $\mathfrak{M}, m_c, a, m_1/h_1 \not\models A$. That is, A is not settled true at m_1 . Sometimes m_1 is called a *counter*.

A moment in $\text{Choice}_w^{Val_\pi(t)}(m)$ is said to be *choice equivalent to m for $Val_\pi(t)$ relative to w* (Def. 12). Thus the positive condition is that A should be settled true at every moment choice equivalent to m for $Val_\pi(t)$ relative to w . In other words, the choice that $Val_\pi(t)$ made at w (where the past tense is from the perspective of m) “guaranteed” that A would be settled true at every moment of $i(m)$ accessible from w via that choice (where the subjunctive is also from the point of view of m , since we are considering alternative ways of filling the same instant).

The negative condition ensures that the choice that $Val_\pi(t)$ made at w was not irrelevant to the truth of A at $i(m)$ in the sense that at w the falsity of A at $i(m)$ is risked—its settled truth there is not already guaranteed at the moment of choice.

Since none of the conditions mention h except in evaluating t as $Val_\pi(t)$, it is clear that $[t \text{ stit}: A]$ is history-independent if t is. Also worth emphasizing is that $[t \text{ stit}: A]$ has a kind of double temporal reference: There is a reference to where A is evaluated and there is a reference to where the witnessing choice occurs. By our semantics, the compound $[t \text{ stit}: A]$ is true (if in fact it is so) where A is evaluated, not at the moment of choice.

We next offer a representation, suitable for use with the achievement stit as in §9B, of the “all in” ability of Austin 1961 (p. 177), present or absent on a particular occasion for a particular agent and with respect to a particular complement. Such ability statements can be tensed either as of the moment of witness or as of the moment of evaluation of the complement. We use “can” for the former and “could have” for the latter:

- $[t_1 \text{ can}_{t_2}\text{-stit: } A]$. READING: Agent t_1 can see to it that at instant t_2 , A . SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models [t_1 \text{ can}_{t_2}\text{-stit: } A]$ iff $\text{Val}_\pi(t_2) \in \text{Instant}$, and there is a moment m_1 in $i|_{>m}$ (i.e., m_1 lies on the horizon from m at i) such that m is a momentary witness to $[t_1 \text{ stit: } A]$ at m_1 relative to $\langle \mathfrak{M}, m_c, a, m/h \rangle$.
- $[t \text{ could-have-stit: } A]$. READING: t could have seen to it that, at the present time, A . SEMANTICS: $\mathfrak{M}, m_c, a, m/h \models [t \text{ could-have-stit: } A]$ iff, where x is a variable not free in $[t \text{ stit: } A]$, $[i_{(m)} / x] \langle \mathfrak{M}, m_c, a, m/h \rangle \models \text{Was:}[t \text{ can}_x\text{-stit: } A]$.

These definitions illustrate the restricted complement thesis, **Thesis 5**. Because $[t \text{ stit: } A]$ occurs as a unit on the right sides of these definitions, you can see that the stit sentence is indeed a complement. You can also see that the right sides would make no sense if $[t \text{ stit: } A]$ were replaced by an arbitrary sentence—which is exactly to say that the complement position is restricted. See also Figure 9.1.

8G.4 Achievement-stit semantics: Witness by chains

A deficiency in the “momentary-witness” concept of stit just defined is that it makes it a matter of “logic” that if Autumn Jane saw to it that she was clean at 4:00, then there was some momentary choice that witnessed that fact. And maybe there was; but suppose instead that the witness to the outcome was a chain of choices by Autumn Jane *with no last member*. Picture her just prior to 4:00 as balancing on a board that crosses over a puddle, and award her the ability to choose to fall off at any of an unending series of moments approaching 4:00. Also permit our illustration to be simple by adding what is probably false: that the lapse from cleanliness occurs as soon as you like after the choice to fall off. Under this supposition the witness to her seeing to her own cleanliness at 4:00 was no single choice, but the whole unending chain of choices properly approaching 4:00. (See the discussion of the ten-minute mile, Question 2-11.)

To keep us all on the right track, we explicitly note that the deficiency to which we point is not at all carried by stories in which in ordinary speech we would say that a long list of preliminary activities went into, say, Autumn Jane’s setting of the table, or to change the illustration in order to make the point even more visible, a story in which in order for Autumn Jane to win at chess at 5:00, a complicated, temporally discrete series of moves was required. The chess illustration makes it clear that catering to these stories has nothing to do with an unending sequence of choices. Reflection on the conceptual problems raised by such stories is important, but these problems are so different from those that we are now discussing that they will require a different treatment. See chapter

13 for some relevant developments via the theory of strategies, and §8G.5 for an additional suggestion.

We turn to the quest for an understanding of the witness of a stit by a possibly unending sequence of choices. We use the notion of choice inseparability explained in **Def. 13**.

- Five conditions must be satisfied by c in order to be a witness to $[t \text{ stit}: A]$ at m relative to $\langle \mathfrak{M}, m_c, a, m/h \rangle$.
 - *Agency.* $Val_\pi(t) \in Agent$.
 - *Priority.* All moments in c must be properly earlier than m .
 - *Nonemptiness.* The chain, c , must of course be nonempty.
 - *Positive condition.* $\mathfrak{M}, m_c, a, m_1/h_1 \models A$ for every moment m_1 and history h_1 such that $m_1 \in i_{(m)} \cap h_1$ and m_1 is inseparable from m for $Val_\pi(t)$ in c (i.e., $m_1 \equiv_c^{Val_\pi(t)} m$; see **Def. 13**).
 - *Negative condition.* Every moment, w , in c “risks” the falsity of A in the sense that above w there is some moment m_1 in $i_{(m)}$ and some history $h_1 \in H_{(m_1)}$ such that $\mathfrak{M}, m_c, a, m_1/h_1 \not\models A$. In other words, for every moment, w , in c , it is not settled at w that A be true at $i_{(m)}$.

The truth conditions for stit, as witnessed by a chain, are as follows.

- $[t \text{ stit}: A]$. **READING:** t sees to it that A . **SEMANTICS:** $\mathfrak{M}, m_c, a, m/h \models A$ iff there is a chain, c , satisfying the agency, priority, nonemptiness, positive, and negative conditions for c to witness $[t \text{ stit}: A]$ at m , relative to $\langle \mathfrak{M}, m_c, a, m/h \rangle$.

See also Definition 13-24 for an equivalent version of the chain-witness semantics for the achievement stit. The idea of chain witness is considered in more detail in Belnap 1996a. Chain witnesses are involved in this book whenever “busy choice sequences” are at issue.

8G.5 The transition stit

Every concept of agency must refer to alternative possibilities, to what might have been without the particular exercise of agency at issue. Which alternatives? Our approach (for better or worse) always seeks an *objective* answer to this question. In the case of dstit, the alternatives are entirely and uniquely determined by the causal structure of the world as encoded by *Tree* and \leq , together with the choices possible for agents as represented by *Choice*. The critical point is that in the case of dstit, the “moment of outcome” (that is, the last moment at which the outcome has not yet begun) is identified with the “moment of choice” (that is, the last moment at which the choice is still hanging in the balance). (See the discussion in §7A.4 of immediate transitions, **Def. 8**.) The alternatives can then be objectively identified as the histories representing

the alternative futures of the moment of choice. When, however, we need to consider outcomes at some temporal remove from the moment or moments of choice, we can no longer be so simple. The achievement-stit solution was to bring in instants, via *Instant*, with which to make “horizontal” comparisons. The alternative possibilities were then objectively identified as alternative ways of filling “the same” instant. Thus in judging whether or not Autumn Jane saw to it that she was clean at 4:00 P.M., we considered what else might have happened at 4:00 P.M. It seems to us likely that there should be other interesting *objective* ways of picking out alternative possibilities.

Here is one, the “transition stit,” that we think has potential for illumination. The first of two fundamental thoughts is that we think about the *transition* aspect of agency, and that it seems to make intuitive sense to say that an agent sees to a transition from a concretely given initial situation or event to a propositionally expressed outcome (see §7A.4). (Both *dstit* and *astit*, in contrast, speak solely of seeing to a propositionally expressed outcome, leaving the initial situation to be brought into play via the moment of evaluation.) The second key thought is that we build into the stit construction itself an explicit reference to a chain of choices or actions witnessing the outcome of the transition. These two thoughts lead us to take as our target, instead of a two-place construction such as *dstit* or *astit*, a four-place construction relating agent, initial event, a given chain of choices, and a propositional outcome.

The linguistic form we carry as $[\alpha \text{ tstit}: m_0 \xrightarrow{c} A]$, the “transition stit,” which we shall think of as being evaluated at a moment m . Because of English tenses, it is difficult to give the transition stit a reading in ordinary language. Permit us to sidestep some of this problem with English by using direct discourse, putting “ A ” in place of the *name* of a sentence, instead of in place of the sentence itself. We may then read $[\alpha \text{ tstit}: m_0 \xrightarrow{c} A]$ as follows.

Given that A was not settled true at m_0 , and considering that the fact that A is settled true (now, at m), it turns out that α is entirely responsible for this transition, which he or she secured by means of a certain chain of choices or actions, namely, c . (13)

The idea of “double time reference” (Definition 5-12 and §8F) hovers in the background. The earlier “initial” moment m_0 is where we evaluate A . But the current moment, m , is where we are evaluating the stit sentence that contains A as a proper part. This is admittedly complicated, but we think that it needs to be so. We should think of A as moment-independent. The reason for this is that we wish to say, in the spirit of double time reference, that at m it is settled true that A , whereas at m_0 it was not settled true that A . In this way we distinguish “causal transition” from simple change of state. For mere change of state, we evaluate A at different moments, as time moves on. On Monday the pond is full, then on Tuesday it is dry. But that sequence is not in its description objectively “causal” since it allows that the pond being dry on Tuesday may have been predetermined from all eternity. In line with §7B.3, the real causal transition needing an explanation is from (something like)

“The pond is dry on Tuesday” is *not* settled true at m_0

to

“The pond is dry on Tuesday” *is* settled true at m_1 .

As is typical in the use of double time references, we *do* evaluate the interior sentence, “The pond is dry,” at different moments. In contrast, the entire sentence, “The pond is dry on Tuesday,” is moment independent. By concentrating on moment-independent sentences, we zero in on the causal transition without getting confused by any change of state. In any event, here is our semantic account of the suggested four-argument “transition stit” representation of seeing to a transition. We include moment independence as a necessary condition.

- $[t_1 \text{ stit}: t_2 \xrightarrow{t_3} A]$. READING: See (13). SEMANTICS: $\mathfrak{M}, m_c, a, m/h$
 $\models [t_1 \text{ stit}: t_2 \xrightarrow{t_3} A]$ iff, letting $\alpha = \text{Val}_\pi(t_1)$, $\alpha \in \text{Agent}$; letting $m_0 = \text{Val}_\pi(t_2)$, $m_0 \in \text{Tree}$; letting $c = \text{Val}_\pi(t_3)$, c is a chain of moments; and
 - *Moment independence of outcome.* $\forall m_1 \forall h_1 [h_1 \in H_{(m_0)} \cap H_{(m_1)} \rightarrow (\mathfrak{M}, m_c, a, m_0/h_1 \models A \Rightarrow \mathfrak{M}, m_c, a, m_1/h_1 \models A)]$;
 - *After and before.* $\forall m_1 [m_1 \in c \rightarrow (m_0 \leq m_1 \text{ and } m_1 < m)]$;
 - *Positive condition.* $\forall h_1 [h_1 \equiv_c^\alpha h \rightarrow \mathfrak{M}, m_c, a, m_0/h_1 \models A]$;
 - *Essentiality condition.* $\forall m_1 [m_1 \in c \rightarrow \exists h_1 [h_1 \in H_{(m_1)} \text{ and } h_1 \perp_{m_1}^\alpha h \text{ and } \mathfrak{M}, m_c, a, m_0/h_1 \not\models A]]$.

Recall, for the positive condition, that $h_1 \equiv_c^\alpha h$ means that h_1 is inseparable from h for α in c , **Def. 13**. The essentiality condition says that no member of c can be omitted. This condition is what gives bite to the idea of the “by means of” phrase in (13). Observe that this is an entirely objective, if limited, concept of “by means of.” The essentiality condition, which says that each member of c is essential, evidently implies a negative condition to the effect that the witnessing is not complete until c is complete.

Even though we do not have the space to develop the idea of stitting a transition by means of a series of choices, we append what seems to us the most promising feature of the idea: that $[\alpha \text{ stit}: m_0 \xrightarrow{c} A]$ encodes agent responsibility for an outcome based on an entire campaign of actions or choices. We are in a position to express, for example, that Mary accomplished the transition from dirty dishes at 7:00 to clean dishes, not by a single action, but by doing one thing after another until the job was done.

This concludes our organized semantic account of some constructions whose exact nature we think important for the understanding of indeterminism. We have also finished the foundational discussions that have interrupted, for the space of three chapters, our consideration of stit and its uses in helping us to understand how agency fits into the causal structure of our world. Next we turn to some applications. Immediately following are two chapters that apply the achievement stit, and then come two involving the deliberative stit.