

Patrick Girard
Olivier Roy
Mathieu Marion
Editors

Dynamic Formal Epistemology

Dynamic Formal Epistemology

SYNTHESE LIBRARY

STUDIES IN EPISTEMOLOGY,
LOGIC, METHODOLOGY, AND PHILOSOPHY OF SCIENCE

Editors-in-Chief:

VINCENT F. HENDRICKS, *University of Copenhagen, Denmark*
JOHN SYMONS, *University of Texas at El Paso, U.S.A.*

Honorary Editor:

JAAKKO HINTIKKA, *Boston University, U.S.A.*

Editors:

DIRK VAN DALEN, *University of Utrecht, The Netherlands*
THEO A.F. KUIPERS, *University of Groningen, The Netherlands*
TEDDY SEIDENFELD, *Carnegie Mellon University, U.S.A.*
PATRICK SUPPES, *Stanford University, California, U.S.A.*
JAN WOLEŃSKI, *Jagiellonian University, Kraków, Poland*

VOLUME 351

For further volumes:
<http://www.springer.com/series/6607>

Dynamic Formal Epistemology

Edited by

Patrick Girard

University of Auckland, New Zealand

Olivier Roy

University of Groningen, The Netherlands

and

Mathieu Marion

Université du Québec à Montréal, Canada



Springer

Editors

Prof. Patrick Girard
Department of Philosophy
University of Auckland
Private Bag 92019
1142 Auckland
New Zealand
p.girard@auckland.ac.nz

Dr. Olivier Roy
University of Groningen
Faculty of Philosophy
Oude Boteringestraat 52
9712 GL Groningen
Netherlands
o.roy@rug.nl

Mathieu Marion
Département de Philosophie
Université du Québec à Montréal
C.P. 8888, Succursale Centre-Ville
Montréal, Québec
Canada H3C 3P8
marion.mathieu@uqam.ca

ISBN 978-94-007-0073-4 e-ISBN 978-94-007-0074-1
DOI 10.1007/978-94-007-0074-1
Springer Dordrecht Heidelberg London New York

© Springer Science+Business Media B.V. 2011

No part of this work may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, microfilming, recording or otherwise, without written permission from the Publisher, with the exception of any material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Acknowledgements

This volume is born out of the workshop *Dynamic Logic Montreal*, held in Montreal in June 2007. As both the editors of this volume and organizers of the workshop, we would like to thank all the participants, attendees, local staff and students, who made this event such a success. The conference had been made possible by a grant from the *Canadian Social Sciences and Humanities Research Council*, which we gratefully acknowledge.

We would like to thank the authors in this volume for their contribution and also for their patience during the lengthy process that led to this publication. We are also grateful to the referees for careful reading and constructive suggestions on each chapter and on the book as a whole.

Finally, we would like to thank Neil Newton, Imogene Paterson, Kym Pram and Vincent Motard-Côté, who helped with the production of this volume, as well as the editors and production team at *Synthese Library*. The editing work has been supported by the University of Auckland's Faculty Research Development Fund.

July 2010

Patrick Girard, Olivier Roy and Mathieu Marion

Contents

1 Introduction	1
Patrick Girard, Mathieu Marion, and Olivier Roy	
2 Logics of Rational Interaction	5
Barteld Kooi and Eric Pacuit	
3 Dynamic Epistemic Logic and Temporal Modality	33
Audrey Yap	
4 Exploring the Power of Converse Events	51
Guillaume Aucher and Andreas Herzig	
5 Modal Logic for Qualitative Dynamics	75
Darko Sarenac	
6 Knowing One's Limits: An Analysis in Centered Dynamic Epistemic Logic	103
Denis Bonnay and Paul Égré	
7 Simple Evidence Elimination in Justification Logic	127
Bryan Renne	
8 Belief Update as Social Choice	151
Johan van Benthem	
9 Revision with Conditional Probability Functions: Two Impossibility Results	161
François Lepage and Charles Morgan	
10 Indeterminacy and Belief Change	173
Horacio Arló-Costa	

11	Perspectival Act Utilitarianism	197
	John F. Horty	
12	Real Change, Deontic Action	223
	Krister Segerberg	
13	Neither Logically Omniscient nor Completely Irrational Agents: Principles for a Fine-Grained Analysis of Propositional Attitudes and Attitude Revision	227
	Daniel Vanderveken	
	Index	239

Contributors

Horacio Arló-Costa Department of Philosophy, Carnegie Mellon University, Baker Hall 135, Pittsburgh, PA 15213-3890, USA, hcosta@andrew.cmu.edu

Guillaume Aucher Faculty of Sciences, Technology and Communication (FSTC), University of Luxembourg, 6 rue Richard Coudenhove – Kalergi L-1359, Luxembourg, guillaume.aucher@uni.lu

Johan van Benthem Institute for Logic, Language & Computation (ILLC), University of Amsterdam, PO Box 94242, 1090 GE, Amsterdam, The Netherlands; Department of Philosophy, Stanford University, Stanford, CA 94305, USA, johan@science.uva.nl; johan.vanbenthem@uva.nl

Denis Bonnay Département de Philosophie, Université Paris Ouest Nanterre, 200 avenue de la République, 92001 Nanterre CEDEX, France; Département d'Etudes Cognitives de l'ENS, IREPH/IHPST, 29 rue d'Ulm, 75005 Paris, France, denis.bonnay@u-paris10.fr; denis.bonnay@ens.fr

Paul Egré Département d'Etudes Cognitives, Institut Jean-Nicod, Ecole Normale Supérieure, 29 rue d'Ulm, 75005 Paris, France, paulegre@gmail.com

Patrick Girard Department of Philosophy, University of Auckland, 18 Symonds Street, Auckland, New Zealand, p.girard@auckland.ac.nz

Andreas Herzig IRIT-LILaC, 118 route de Narbonne, F-31062, Toulouse Cedex 9, France, herzig@irit.fr

John F. Horty Philosophy Department, Institute of Advanced Computer Studies, University of Maryland, College Park, MD 20742, USA, horty@umiacs.umd.edu

Barteld Kooi Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands, B.P.Kooi@rug.nl

François Lepage Département de Philosophie, Université du Québec à Montréal, C.P. 6128, succursale Centre-ville, Montréal, QC H3C 3J7, Canada, francois.lepage@umontreal.ca

Mathieu Marion Département de Philosophie, Université du Québec à Montréal, C.P. 8888, Succursale Centre-Ville, Montréal, Québec, Canada H3C 3P8, marion.mathieu@uqam.ca

Charles Morgan Department of Mathematics, University College London, Gower Street, London WC1E 6BT, UK, charles.morgan@ucl.ac.uk

Eric Pacuit Center for Logic and Philosophy of Science, Tilburg University, 5000 LE Tilburg, The Netherlands, e.j.pacuit@uvt.nl

Bryan Renne Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands, bryan@renne.org

Olivier Roy Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands, o.roy@rug.nl

Darko Sarenac Department of Philosophy, Colorado State University, Fort Collins, CO 80523-1781, USA, darko@colostate.edu

Krister Segerberg Filosofiska institutionen, Uppsala Universitet, 751 26 Uppsala, Sweden, krister.segerberg@filosofi.uu.se

Daniel Vanderveken Département de Philosophie, Université du Québec à Trois-Rivières, Trois-Rivières, QC G9A 5H7, Canada, daniel.vanderveken@uqtr.ca

Audrey Yap Department of Philosophy, University of Victoria, Victoria, BC V8W 3P4, Canada, ayap@uvic.ca

Chapter 1

Introduction

Patrick Girard, Mathieu Marion, and Olivier Roy

The frontier of contemporary epistemology is dynamic. Shifting from purely conceptual analysis, the theory of individual knowledge, belief and justification now includes an increasing amount of formal work, utilizing either logic or probabilities. Epistemology has also moved to questions regarding information change, its flow among groups, and its place within interaction, whereas for a long time it was centered mainly on the question of individual knowledge and its acquisition in a static environment. Epistemology is thus expanding beyond the conceptual analysis of justified true belief, allowing for a broad and formal philosophical inquiry into the notion of information and how it is acquired, changed, passed on, and aggregated. By doing so it provides new insights and methods relevant not only to the theory of knowledge but also for our understanding of interaction, obligations, and scientific discovery.

Dynamic epistemology played a part in the renewal of formal analytical philosophy. Numerous seminal contributions in the middle of the twentieth century were closely connected with formal investigations: A. Prior on time and determinism, G. H. von Wright on preferences and obligations, S. Kripke on direct reference, J. Hintikka on knowledge and beliefs, D. Lewis on conditionals and conventions, and R. Jeffrey on Bayesian rationality, to name but a few. Since the 1970s, however, formal work within the field of logic has moved slowly away from analytic philosophy towards mathematics, computer science and linguistics. These different areas have provided logic with a plethora of new mathematical tools, and have aided the development of techniques that can analyze information flow, belief revision, preference change, and strategic interaction. These tools proved to be much more than just fruitful applications; philosophers now realize that they also shed new light on foundational questions.

This book brings together original contributions from the actors of this dynamic turn in epistemology. It aims to bring their work under a single umbrella by highlighting the coherence of their current research themes, and by establishing connec-

O. Roy (✉)

Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, 9712 GL Groningen, The Netherlands

e-mail: o.roy@rug.nl

tions between topics that up until now have been investigated independently. It will also be a helpful red for any analytic philosopher who is not yet acquainted with the dynamic perspective, as it illustrates how the new analytical toolbox unveils fresh questions about the theory of knowledge, belief, preference, action, and rationality.

The contributions in this book explore a number of central axes in dynamic epistemology: temporal, social, probabilistic and even deontic dynamics. This diversity is by no means a sign of disunity; rather, the dynamic way of thinking sheds light on a broad array of topics. The following is a short overview of these contributions.

Temporal information change is an obvious subject to arise from the study of dynamic epistemology. An interest in time and events is certainly not new in philosophy, especially in the analytic or the logical tradition, but the emergence of propositional dynamic logic (PDL), epistemic temporal logic (ETL) and dynamic epistemic logic (DEL) has revealed a number of crucial issues involving temporal reasoning. In [Chapter 2](#), Eric Pacuit and Barteld Kooi provide the unaccustomed reader with an introduction to these formal frameworks. For the expert they demonstrate that PDL, ETL and DEL are not competing systems; their complementarity supersedes their differences. The reader will also find a comprehensive bibliography of the themes covered in this book.

Dynamic epistemic logic, in its standard form, suffers from a widespread limitation, which [Chapters 3](#) and [4](#) endeavor to tackle. In their most common forms, these logical languages only refer to current information states and their transformations, and do not refer back to how things were before certain events occurred. In [Chapter 3](#), Audrey Yap avoids this limitation by extending the standard dynamic epistemic language to include a past operator. She provides a sound and complete axiomatization for this new inclusion and investigates its expressive power. In [Chapter 4](#), Guillaume Aucher and Andreas Herzig investigate past operators in propositional dynamic logic. Both contributions illustrate the interconnections, demonstrated by Pacuit and Kooi in [Chapter 2](#), between the different logical systems used to talk about temporal information change. Yaps models for DEL are similar to ETL structures, while Aucher and Herzig show that standard DEL can be faithfully embedded into their extended PDL.

In [Chapter 5](#), Sarenac investigates a general approach to dynamic systems, using a notion borrowed from computer science, namely iterative function systems (IFS). Dynamic epistemic logic is one among many categories of dynamic problems that can be analyzed in this setting. Indeed, Sarenac traces his analysis back to Poincaré's work on the three-body problem, and contrasts dynamical analyses in mathematical physics with those more common in computer science, which are central to the present book.

Dennis Bonnay and Paul Egge show in [Chapter 6](#) that, when one considers temporal dynamics, elegant solutions may be found to questions regarding imprecise knowledge. They provide a dynamic analysis of Timothy Williamson's Margin of Error Paradox. The analysis shows that Williamsons paradox stems from a rather simplistic assumption about the rigidity of margins of error through time.

In [Chapter 7](#), Bryan Renne addresses the topic of evidence, and in particular, evidence elimination, which is another issue involved in the study of temporal change.

Using justification logic, a logic that can analyze the evidence put forward to support a proof, Renne provides a sound and complete axiomatic system for the logic of evidence change in a way that, again, bridges different logical systems this time between justification and dynamic epistemic logics.

In [Chapter 8](#), Johan van Benthem bridges the gap between the single-agent perspective and the dynamics of social interaction. He shows that well-known dynamic epistemic logical systems, originally designed to analyze information updates after epistemic events, can be seen as systems of preference aggregation commonly used in Social Choice Theory. He does so by providing a characterization of the standard DEL update rule in terms of a priority update, which opens up a whole new perspective on dynamic epistemic logical systems.

The study of belief change already has a long history within Bayesian and probabilistic approaches to epistemology. Despite this, the following two contributions demonstrate that if we take information dynamics seriously, there are still major challenges that the existing approaches face.

In [Chapter 9](#), François Lepage and Charles Morgan take Lewis well-known triviality result for the conditional probability of a counterfactual and generalize it to any two-place probability function satisfying minimal requirements. This result applies to a wide range of belief change operations from classical conditioning to imaging. Alternative probabilistic views on counterfactual, reasoning and belief change are called for if one is to pursue this approach.

In [Chapter 10](#), Horacio Arló Costa puts forward another important challenge to contemporary probabilistic approaches to belief change, namely, how one may account for such phenomenon in situations where attitudes are incomplete or indeterminate. Arló Costa sketches a two-tier theory of belief change and presents the various challenges that its formalization poses. He also considers the application of this theory to Philosophy of Science, thus illustrating how dynamic epistemology can contribute to traditional debates about the theory of knowledge and scientific inquiry.

The next two chapters in this collection show that notions of obligation can be analyzed dynamically, thus extending the scope of dynamic studies to the deontic realm. John Horty's paper in [Chapter 11](#), which nicely complements the temporal logic of the first chapters, looks at deontic modalities in branching time structures. In previous works Horty has demonstrated that an elegant formalization of act utilitarianism can be made by extending Belnap's stit (see to it that) logic with deontic operators. In this contribution, he argues that various interpretations of these operators are possible in stit models, interpretations that seem equally plausible but which turn out to make contradicting assessments of actions in certain contexts. He shows, however, that by borrowing the idea of double time references from temporal logic, a very general account of act utilitarianism based on different perspectives can be given.

In [Chapter 12](#), Krister Segerberg addresses issues pertaining to deontic logic over temporal structures. He proposes a formalism to deal with factual and normative changes, and the interaction between the two, thus improving on some of his earlier work.

Finally, in [Chapter 13](#), Daniel Vanderveken proposes a general logical typology of propositional attitudes, ranging from the basic notions of belief, desire and intention, to sophisticated notions such as regret and expectation. In his chapter Vanderveken shows that his approach avoids the pitfalls of logical omniscience by providing a fine-grained intensional semantic. He also provides a general theory of attitude revision that can deal with all kinds of attitude changes.

* * * *

The contributions in this volume result from the workshop *Dynamic Logic Montreal*, held at the Université du Québec à Montréal (UQAM) in June 2007. We would like to thank the Social Sciences and Humanities Research Council of Canada (SSHRC), UQAM, the Université de Montréal, the Canada Research Chair in the Philosophy of Logic and Mathematics and the Institute for Logic, Language and Computation for their financial support. We are also grateful to: Imogene Paterson, Neil Newton, Kym Pram and Vincent Motard-Côté for their help during the workshop and/or during the preparation of this volume; the referees of the various chapters, and of the volume as a whole, for their useful comments and suggestions; and Vincent Hendricks Ingrid van Laarhoven at Synthese Library for their enthusiasm and support. Finally, thanks to all contributors of this book for their efforts, and for their patience, during the rather lengthy process that lead to this publication.

Chapter 2

Logics of Rational Interaction

Barteld Kooi and Eric Pacuit

2.1 Introduction

There is a growing literature focused on using logical methods to reason about communities of agents engaged in some form of social interaction. Much of the work builds upon existing logical frameworks developed by philosophers and computer scientists incorporating insights and ideas from philosophy (especially epistemology and philosophy of action), game theory, decision theory and social choice theory. The result is a web of logical systems each addressing different aspects of rational agency and social interaction. Rather than providing an encyclopedic account of these different logical systems,¹ this chapter focuses on issues surrounding the modeling of informational attitudes in social interactive situations. The main objective is to introduce the two main approaches to modeling “rational interaction” and provide pointers to the current literature.

Of course, there is no single approach that can address *all* of the complex phenomena that arise when rational agents interact with one another and the environment. Thus it is important to understand how the different analyses from within and across the disciplines mentioned above can fit together. This suggests the following three general questions:

1. How can we *compare* different logical frameworks addressing similar aspects of rational agency and social interaction (i.e., how information evolves through social interaction)?
2. How should we *combine* logical systems which address *different* aspects of social interaction towards the goal of a comprehensive (formal) theory of rational agency?

B. Kooi (✉)

Faculty of Philosophy, University of Groningen, Oude Boteringestraat 52, 9712 GL, Groningen, The Netherlands
e-mail: B.P.Kooi@rug.nl

¹The interested reader can consult Meyer and Veltman (2007), van der Hoek and Wooldridge (2003), van Benthem (2008a) and references therein.

3. How do the logical frameworks discussed in this literature contribute to the broader discussion of rational agency and social interaction within philosophy and the social sciences?

Certainly, the first two questions raise numerous methodological issues and technical problems. However, they also make explicit certain foundational and philosophical issues surrounding rational interaction (cf. van Benthem et al. 2009a). In particular, viewing the various logical systems found in the literature as (sometimes competing) accounts of rational agency forces us to carefully examine what we even mean by a “rational agent” (see van Benthem 2005, for an extensive discussion). Of course, the nature of rationality and human agency is a central concern of many philosophers from Aristotle to David Hume to present-day philosophers (cf. Bratman 2007, Searle 1985, Hyman and Steward 2004). The point here is that there are many different types of reasoning and dynamic processes that agents use when interacting with other agents. Comparing and combining the different logical systems forces us to consider how these different processes interact.

In this survey, the modeling of informational attitudes of a group of *rational* agents engaged in some form of social interaction (e.g. having a conversation or playing a card game) takes center stage. Indeed, there are many logical systems today that describe how an agent’s information changes over time. Sometimes the differences between two competing logical systems are technical in nature reflecting different conventions used by different research communities. And so, with a certain amount of technical work, such frameworks are seen to be equivalent up to model transformations (cf. Halpern 1999, Lomuscio and Ryan 1997, Pacuit 2007a, van Benthem et al. 2009a). Other differences point to key conceptual issues about rational interaction. We will introduce the two main logical accounts of rational interaction and highlight such similarities and differences.

2.2 Reasoning About Rational Interaction

This section introduces two logical frameworks that describe the dynamics of information over time in a multiagent situation. The first is *epistemic temporal logic* (ETL, Fagin et al. 1995, Parikh and Ramanujam 1985) which uses linear or branching time models with added epistemic structure induced by the agents’ different capabilities for observing events. These models provide a “grand stage” where histories of some social interaction unfold constrained by a **protocol**. Here a **protocol** is intended to represent the rules or conventions that govern many of our social interactions. For example, in a conversation, it is typically not polite to “blurt everything out at the beginning”, as we must speak in small chunks. Other natural conversational protocol rules include “do not repeat yourself”, “let others speak in turn”, and “be honest”. Imposing such rules *restricts* the legitimate sequences of possible statements.

The other framework is *dynamic epistemic logic* (DEL, Gerbrandy 1999a; Baltag et al. 1998b; van Ditmarsch et al. 2007) that describes social interactions in terms of epistemic **event models** (which may occur inside modalities of the language). Similar to the way Kripke structures are used to capture the information the agents have about a *fixed* social situation,² an **event model** describes the agents' information about which actual events are currently taking place. The temporal evolution of the situation is then computed from some initial epistemic model through a process of successive "product updates". Details of both frameworks are provided in the subsequent sections.

Often DEL and ETL are presented as *competing* ways of adding dynamics to multi-agent epistemic models. Based on van Benthem et al. (2009, 2006) and van Benthem and Pacuit (2006), we will see how DEL and ETL should rather be viewed as *complementary* accounts of social interaction. The focus is on conceptual issues leaving some of the more technical details and proofs to the relevant papers. The following running example will help guide intuitions (also discussed in Pacuit and Parikh 2006).

Example 2.1 Suppose that Ann would like Bob to attend her talk; however, she only wants Bob to attend if he is interested in the subject of her talk, not because he is just being polite. There is a very simple procedure to solve Ann's problem: Have a (trusted) friend tell Bob the time and subject of her talk.

Taking a cue from computer science, perhaps we can *prove* that this simple procedure correctly solves Ann's problem. However, it is not so clear how to define a correct solution to Ann's problem. If Bob is actually present during Ann's talk, can we conclude that Ann's procedure succeeded? Not really. Bob may have figured out that Ann wanted him to attend, and so is there only out of politeness. Thus for Ann's procedure to succeed, she must achieve a certain "level of knowledge" (cf. Parikh, 2003) between her and Bob. Besides both Ann and Bob knowing about the talk and Ann knowing that Bob knows about

Bob *does not know* that Ann knows about the talk.

This last point is important, since, if Bob knows that Ann knows that he knows about the talk, he may feel social pressure to attend.³ Thus, the procedure *to have a friend tell Bob about the talk, but not reveal that it is at Ann's suggestion*, will satisfy all the conditions. Telling Bob directly will satisfy the first three, but not the essential last condition.

²A Kripke structure is a set of states with relations on this set for each agent. The states, or possible worlds, represent different ways the social situation could have evolved and the relations describe the agents' (current) information. See, for example, Fagin et al. (1995) for details.

³Of course, this is not meant to be a complete analysis of "social politeness".

2.2.1 Epistemic Temporal Logic

Fix a finite set of agents $\mathcal{A}$ and a (possibly infinite) set of events⁴ Σ . A **history** is a finite sequence of events⁵ from Σ . We write Σ^* for the set of histories built from elements of Σ . For a history h , we write he for the history h followed by the event e . Given $h, h' \in \Sigma^*$, we write $h \preceq h'$ if h is a prefix of h' , and $h \prec_e h'$ if $h' = he$ for some event e .

For example, consider the social interaction described in Example 2.1. There are three relevant participants: Ann (A), Bob (B) and Ann's friend (call him Charles (C)). What are the relevant primitive events? To keep things simple, assume that Ann's talk is either at 2PM or 3PM and initially none of the agents know this. Say, that Ann receives a message stating that her talk is at 2PM (denote this event – Ann receiving a private message saying that her talk is at 2PM – by e_A^{2PM}). Now, after Ann receives the message that the talk is at 2PM, she proceeds to tell her trusted friend Charles that the talk is at 2PM (and that she wants him to inform Bob of the time of the talk without acknowledging that the information can from her – call this event e_C^A), then Charles tells Bob this information (call this event e_B^C). Thus, the history

$$e_A^{2PM} e_C^A e_B^C$$

represents the sequence of events where “Ann receives a (private) message stating that the talk is at 2PM, Ann tells Charles the talk is at 2PM, then Charles tells Bob the talk is at 2PM”. Of course, there are other events that are also relevant to this situation. For one thing, Ann could have received a message stating that her talk is at 3PM (denote this event by e_A^{3PM}). This will be important to capture Bob's uncertainty about whether Ann knows that he knows about the talk. Furthermore, Charles may learn about the time of the talk independently of Ann (denote these two events by e_C^{2PM}, e_C^{3PM}). So, for example, the history

$$e_A^{2PM} e_C^{2PM} e_B^C$$

represents the situation where Charles independently learns about the time of the talk and informs Bob.

There are a number of simplifying assumptions that we adopt in this section. They are not crucial for the analysis of Example 2.1, but do simplify the some of

⁴There is a large literature addressing the many subtleties surrounding the very notion of an *event* and when one event *causes* another event (see, for example, Cartwright 2007). However, for this chapter we take the notion of event as primitive. What is needed is that if an event takes place at some time t , then the fact that the event took place can be observed by a relevant set of agents at t . Compare this with the notion of an event from probability theory. If we assume that at each clock tick a coin is flipped exactly once, then “the coin landed heads” is a possible event. However, “the coin landed heads more than tails” would not be an event, since it cannot be observed at any one moment. As we will see, the second statement will be considered a *property* of histories, or sequences of events.

⁵To be precise, elements of Σ should, perhaps, be thought of as event *types* whereas elements of a history are event *tokens*.

the formal details. Since, histories are sequences of (discrete) events, we assume the existence of a global discrete clock (whether the agents have access to this clock is another issue that will be discussed shortly). The length of the history then represents the amount of time that has passed. Note that this implies that we are assuming a finite past with a possibly infinite future. Furthermore, we assume that at each clock tick, or moment, *some* event takes place (which need not be an event that any agent directly observes). Thus, we can include an event e_t (for “clock tick”) which can represent that “Charles does *not* tell Bob that the talk is at 2PM.” So the history


$$e_A^{2PM} e_C^A e_t$$

describes the sequence of events where, after learning about the time of the talk, Ann informs Charles, but Charles does *not* go on to tell Bob that the talk is at 2PM. Once a set of events Σ is fixed, the temporal evolution and moment-by-moment uncertainty of the agents can be described.

Definition 2.1 (ETL Frames) Let Σ be a set of events. A **protocol** is a set $H \subseteq \Sigma^*$ closed under non-empty prefixes. An **ETL frame** is a tuple $\langle \Sigma, H, \{\sim_i\}_{i \in \mathcal{A}} \rangle$ with H a protocol, and for each $i \in \mathcal{A}$, a binary relation $\sim_i$ on⁶ H .

An ETL frame describes how the agents’ *hard* information⁷ evolves over time in some social situation. The protocol describes (among other things) the temporal structure, with h' such that $h <_e h'$ representing the point in time after e has happened in h . The relations $\sim_i$ represent the uncertainty of the agents about how the current history has evolved. Thus, $h \sim_i h'$ means that from agent i ’s point of view, the history h' looks the same as the history h .

Note that the protocol in an ETL frame captures not only the temporal structure of the social situation being modeled but also assumptions about the nature of the participants. For example, the following is a possible protocol built from the events described above:



⁶Although we will not do so here, typically it is assumed that $\sim_i$ is an equivalence relation.

⁷As opposed to *soft* information which may be revised. See van Benthem (2007) for a general discussion of hard and soft information.


While this protocol does describe possible ways the situation described in Example 2.1 could evolve, it does not account for the *motivation* of the agents. For example, the history

$$e_A^{3PM} e_C^A e_B^C$$

describes the sequence of events where Ann learns the talk is at 3PM but tells Charles (who goes on to inform Bob) that the talk is at 2PM. Of course, given the assumption that Ann *wants* Bob to attend her talk, this should not be part of (Ann's) protocol. Similarly, since we assume Charles is trustworthy, we should not include any histories where e_t follows the event e_C^A . Taking into account these underlying assumptions about the motivations (e.g. Ann wants Bob to attend the talk) and dispositions (e.g. Charles tells the truth and lives up to his promises) of the agents we can drop a number of histories from the protocol shown above. Note that we keep the history

$$e_A^{2PM} e_C^{2PM} e_t$$

in the protocol, since if Charles learns independently about the time of the talk, then he is under no obligation to inform Bob. In the picture below, we also add some of the uncertainty relations for Ann and Bob (to keep the picture simple, we do not draw the full ETL frame). The solid line represents Bob's uncertainty while the dashed line represents Ann's uncertainty. The main assumption is that Bob can only observe the event (e_B^C). So, for example, the histories $h = e_A^{2PM} e_C^A e_B^C$ and $h' = e_A^{2PM} e_C^{2PM} e_B^C$ look the same to Bob (i.e., $h \sim_B h'$).⁸



Assumptions about the underlying protocol in an ETL frame corresponds to “fixing the playground” where the agents will interact. As we have seen, the protocol not only describes the temporal structure of the situation being modeled, but also any *causal* relationships between events (e.g., sending a message must always proceed receiving that message) plus the motivations and dispositions of the participants (e.g., liars send messages that they *know* – or believe – to be false). Thus

⁸Note that we do not include any reflexive arrows in the picture in order to keep things simple.

the “knowledge” of agent i at a history h in some ETL frame is derived from both i ’s observational powers (via the $\sim_i$ relation) and i ’s information about the (fixed) protocol.

Remark 2.1 (Three Equivalent Approaches) There are at least two further approaches to uncertainty in the literature. The first, discussed by Parikh and Ramanujam (1985), explicitly describes the agents’ “observational” power. That is, each agent i has a set E_i of events she *can* observe.⁹ For simplicity, we assume $E_i \subseteq \Sigma$ but this is not necessary. A **local view** function is a map $\lambda_i : \mathbf{H} \rightarrow E_i^*$. Given a finite history $h \in \mathbf{H}$, the intended interpretation of $\lambda_i(h)$ is “the sequence of events observed by agent i at h ”. The second approach comes from Fagin et al. (1995). Each agent has a set L_i of **local states** (if necessary, one can also assume a set L_e of environment states). Events e are tuples of local states (one for each agent) $\langle l_1, \dots, l_n \rangle$ where for each $i = 1, \dots, n$, $l_i \in L_i$. Then two finite histories h and h' are i -equivalent provided the local state of the last of event on h and h' is the same for agent i . From a technical point of view, the three approaches (uncertainty relations, local view functions and local states) to modeling uncertainty are equivalent (Pacuit 2007a, van Benthem and Pacuit 2006, provide the relevant discussions).

Although, syntactic issues do not play an important role in this chapter, we give the bare necessities to facilitate a comparison between ETL and DEL. Different modal languages describe ETL frames (see, for example, Hodkinson and Reynolds 2006, Fagin et al. 1995), with “branching” or “linear” variants. Let $\mathbf{At}$ be a countable set of atomic propositions. The language $\mathcal{L}_{ETL}$ is generated by the following grammar:

$$P \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi \mid \langle e \rangle\varphi$$

where $i \in \mathcal{A}$, $e \in \Sigma$ and $P \in \mathbf{At}$. The usual boolean connectives ($\vee$, $\rightarrow$, $\leftrightarrow$) and the dual modal operators (L_i , $[e]$) are defined as usual. The pure epistemic language, denoted $\mathcal{L}_{EL}$, is the fragment of $\mathcal{L}_{ETL}$ with only epistemic modalities (which we will refer to both as the “language of epistemic logic” and the “epistemic fragment” of $\mathcal{L}_{ETL}$ or the language $\mathcal{L}_{DEL}$ defined below). The intended interpretation of “ $K_i\varphi$ ” is “according to agent i ’s current information, φ is true.” The intended interpretation of “ $\langle e \rangle\varphi$ ” is “after event e (does) take place, φ is true.” Formulas are interpreted at histories in an **ETL model**:

Definition 2.2 (ETL Model) An **ETL model** is a tuple $\langle \Sigma, \mathbf{H}, \{\sim_i\}_{i \in \mathcal{A}}, V \rangle$ with $\langle \Sigma, \mathbf{H}, \{\sim_i\}_{i \in \mathcal{A}} \rangle$ an ETL frame and V a valuation function ($V: \mathbf{At} \rightarrow 2^{\mathbf{H}}$).

Definition 2.3 (Truth of $\mathcal{L}_{ETL}$ Formulas) Let $\mathcal{H} = \langle \Sigma, \mathbf{H}, \{\sim_i\}_{i \in \mathcal{A}}, V \rangle$ be an ETL model. The truth of a formula φ at a history $h \in \mathbf{H}$, denoted $\mathcal{H}, h \models \varphi$, is defined inductively as follows:

1. $\mathcal{H}, h \models P$ iff $h \in V(P)$
2. $\mathcal{H}, h \models \neg\varphi$ iff $\mathcal{H}, h \not\models \varphi$
3. $\mathcal{H}, h \models \varphi \wedge \psi$ iff $\mathcal{H}, h \models \varphi$ and $\mathcal{H}, h \models \psi$

⁹This may be different from what the agent *does* observe in a given situation.

4. $\mathcal{H}, h \models K_i \varphi$ iff for each $h' \in \mathbf{H}$, if $h \sim_i h'$ then $\mathcal{H}, h' \models \varphi$
5. $\mathcal{H}, h \models \langle e \rangle \varphi$ iff there exists an $h' \in \mathbf{H}$ such that $h \prec_e h'$ and $\mathcal{H}, h' \models \varphi$

It is often natural to extend the language $\mathcal{L}_{ETL}$ with group knowledge operators (e.g., common or distributed knowledge) and more expressive temporal operators (e.g., arbitrary future or past modalities).

2.2.2 Dynamic Epistemic Logic


An alternative account of interactive dynamics was elaborated by Gerbrandy (1999a), Baltag et al. (1998b), van Benthem (2006), van Benthem et al. (2006) and others. From an initial epistemic model, temporal structure evolves as explicitly triggered by complex informative events.

Definition 2.4 (Epistemic Model) Let $\mathcal{A}$ be a finite set of agents and $\mathbf{At}$ a set of atomic propositions. An **epistemic model** is a tuple $\langle W, \{R_i\}_{i \in \mathcal{A}}, V \rangle$ where W is a non-empty set, for each $i \in \mathcal{A}$, R_i is a relation¹⁰ on W ($R_i \subseteq W \times W$) and V a valuation function ($V : \mathbf{At} \rightarrow 2^W$). We call the set W the domain of $\mathcal{M}$, denoted by $D(\mathcal{M})$. A pair $\mathcal{M}, w$ where $\mathcal{M}$ is an epistemic model and $w \in D(\mathcal{M})$ is called a **pointed epistemic model**.

We can interpret the epistemic language, $\mathcal{L}_{EL}$, defined above at states in an epistemic model. Truth is defined as usual. We only recall the definition of the knowledge operators:

$$\mathcal{M}, w \models K_i \varphi \text{ iff for each } w' \in W, \text{ if } w R_i w' \text{ then } \mathcal{M}, w' \models \varphi$$

Returning to our running example (Example 2.1), initially we assume that none of the agents knows the time of Ann’s talk. Let P be the proposition “Ann’s talk is at 2PM.” Then this initial model can be pictured as follows: there are two states w and v with P true at w ($w \in V(P)$). The agent’s uncertainty relations is the universal relation (since all agents have the same information, we do not label the arrows). Note that the convention followed in this section is that a solid line around a state means that state is the *actual* or current state (i.e., where the formulas are to be evaluated):




Whereas an ETL frame describes the agents’ information at all moments, **event models** are used to build new epistemic models as needed.

Definition 2.5 (Event Model) An **event model** is a tuple $\langle S, \{\longrightarrow_i\}_{i \in \mathcal{A}}, \text{pre} \rangle$, where S is a nonempty set of **primitive events**, for each $i \in \mathcal{A}$, $\longrightarrow_i \subseteq S \times S$


¹⁰Again, the R_i are often taken to be equivalence relations on W – but we do not commit.

and $\text{pre } S \rightarrow \mathcal{L}_{EL}$ is the **pre-condition function**. The set S in an event model $\mathcal{E}$ is called the domain of $\mathcal{E}$, denoted $D(\mathcal{E})$.

Given two primitive events e and f , $e \rightarrow_i f$ means that “according to agent i , event e looks like event f .” Event models then describe an “epistemic event”. In Example 2.1 the first event is Ann receiving a private message that the talk is at 2PM. This can be described by a simple event model: there are two primitive events e and f . The precondition of e is P ($\text{pre}(e) = P$) and the precondition of f is $\top$ (i.e., f is the “skip event”).



Thus, initially Ann observes the actual event e (and so, learning that P is true) while Bob and Charles observe a skip event (and so, their information does not change). What is the effect of this event on the initial model pictured above? Intuitively, it is not hard to see that after the initial event, Ann knows that P is true while Bob and Charles are still ignorant of P and the fact that Ann knows P . That is, combining the initial epistemic model with the above event model should yield the following epistemic model (for simplicity we only draw Ann and Bob’s uncertainty relations):




The following definition gives a general procedure for constructing a new epistemic model from a given epistemic model and an event model.


Definition 2.6 (Product Update) The **product update** $\mathcal{M} \otimes \mathcal{E}$ of an epistemic model $\mathcal{M} = \langle W, \{R_i\}_{i \in \mathcal{A}}, V \rangle$ and event model $\mathcal{E} = \langle S, \{\rightarrow_i\}_{i \in \mathcal{A}}, \text{pre} \rangle$ is the epistemic model $\langle W', R'_i, V' \rangle$ with

1. $W' = \{(w, e) \mid w \in W, e \in S \text{ and } \mathcal{M}, w \models \text{pre}(e)\}$,
2. $(w, e)R'_i(w', e') \text{ iff } wR_iw' \text{ in } \mathcal{M} \text{ and } e \rightarrow_i e' \text{ in } \mathcal{E}, \text{ and}$
3. For all $P \in \text{At}$, $(s, e) \in V'(P) \text{ iff } s \in V(P)$

We illustrate this construction using our running example. The main event in Example 2.1 is “Charles telling Bob (without Ann present) that Ann’s talk is at 2PM”. This can be described using the following event model (again only the Ann and Bob relations will be drawn): Ann is aware of the actual event taking place while Bob thinks the event is a private message to himself.



As in the previous section, there are implicit assumptions here about the motivations and dispositions of the agents. Thus, even though Ann is not present during the actual event,¹¹ she *trusts* that Charles will honestly tell Bob that the talk is at 2PM (without revealing he received the information from her). This explains why in the above event model, $e_1 \xrightarrow{A} e_1$. Starting from a slightly modified epistemic model from the one given above (where Bob now knows that Ann knows *whether* the talk is at 2PM), using Definition 2.6, we can calculate the effect of the above event model as follows (again focusing only on Ann and Bob’s information):



¹¹ Of course, we must assume that she knows precisely *when* Charles will meet with Bob.

Note that, in the epistemic model on the right, for simplicity, the reflexive arrows are not drawn.

Finally, a few comments about syntactic issues. The language $\mathcal{L}_{\text{DEL}}$ extends $\mathcal{L}_{\text{EL}}$ with operators $\langle \mathcal{E}, e \rangle$ for each pair of event models $\mathcal{E}$ and event e in the domain of $\mathcal{E}$. Truth for $\mathcal{L}_{\text{DEL}}$ is defined as usual. We only define the typical DEL modalities:

$$\mathcal{M}, w \models \langle \mathcal{E}, e \rangle \varphi \text{ iff } \mathcal{M}, w \models \text{pre}(e) \text{ and } \mathcal{M} \otimes \mathcal{E}, (w, e) \models \varphi$$

Example 2.2 Public Announcement Logic The **public announcement** of a formula $\varphi \in \mathcal{L}_{\text{EL}}$ is the event model $\mathcal{E}_\varphi = \langle \{e\}, \{\longrightarrow_i\}_{i \in \mathcal{A}}, \text{pre} \rangle$ where for each $i \in \mathcal{A}$, $e \longrightarrow_i e$ and $\text{pre}(e) = \varphi$ (see Plaza 2007, Gerbrandy 1999a). As the reader is invited to verify, the product update of an epistemic model $\mathcal{M}$ with a public announcement model $\mathcal{E}_\varphi$ is the submodel of $\mathcal{M}$ containing all the states that satisfy φ . In this case, the DEL modality $\langle \mathcal{E}_\varphi, e \rangle$ will be denoted $\langle \varphi \rangle$. Henceforth, $\mathcal{L}_{\text{PAL}}$ will denote this language.

2.2.3 Comparing DEL and ETL

Both ETL and DEL are logical frameworks that describe the flow of information in a social interactive situation. For instance the *broadcasts* studied by van der Meyden (1996) and Lomuscio et al. (2000) are essentially the *public announcements* of Example 2.2. So, it is natural to ask how these two frameworks are related (cf. question 1 from the Introduction). Different logical frameworks, such as DEL and ETL, can be compared along many different dimensions. One key way to compare two different logical frameworks focuses on their *expressivity*. In order to show that one logic is at least as expressive as another logic, there are two main tasks to be carried out:

1. One has to establish a relation between the models of the two logics so that if we are given a model from the one logic, we can construct a corresponding model for the other logic;
2. One has to provide a formal translation so that if we are given a formula in the one formal language, we can produce a formula in the other with the same meaning.

Connections¹² between DEL and ETL along these lines have been worked out in detail by van Benthem and Pacuit (2006) and van Benthem et al. (2009).

The key observation is that by repeatedly updating an epistemic model with event models, the machinery of DEL (i.e., Definition 2.6) in effect creates ETL models. Note that an ETL model contains not only a description of how the agents' information changes over time, but also "protocol information" describing *when* each

¹²The first formal connection was established by Gerbrandy (1999a, Section 5.3).